

Spectrum-based Software Fault Localization: A Survey of Techniques, Advances, and Challenges

Higor A. de Souza^{*1}, Marcos L. Chaim^{†2}, and Fabio Kon^{‡1}

¹Department of Computer Science – University of São Paulo

²School of Arts, Sciences, and Humanities – University of São Paulo

Abstract

Despite being one of the most basic tasks in software development, debugging is still performed in a mostly manual way, leading to high cost and low performance. To address this problem, researchers have studied promising approaches, such as Spectrum-based Fault Localization (SFL) techniques, which pinpoint program elements more likely to contain faults. This survey discusses the state-of-the-art of SFL techniques, including their costs, the type and number of faults they address, the programs they utilize in their validation, the testing information that support them, and their use at industrial settings. Notwithstanding the advances, there are still challenges for the industry to adopt these techniques, which we analyze in this paper. This includes using program analysis to generate reduced sets of suspicious entities; combining different spectra to fine-tune the fault localization ability; using strategies to collect fine-grained coverage levels from suspicious coarser levels for balancing execution costs and output precision; and proposing new techniques to cope with multiple-fault programs. Moreover, additional user studies are needed to better understand how SFL techniques can be used in practice. We conclude by presenting a concept map about topics and challenges for future research in SFL.

Keywords: Fault localization, spectrum-based, coverage-based.

1 Introduction

Program faults are an inevitable consequence of writing code. Faults occur for various reasons: typing errors, misunderstanding software requirements, wrong values assigned to variables, or absence of code to verify some unpredicted condition. During the testing phase or in the field, a fault is generally revealed when a program presents an unexpected behavior (known as failure). Once a failure occurs, the two-step debugging process begins. First, a developer inspects the code to *locate* the failure’s cause. Second, s/he *fixes* the fault [Myers, 1979].

Fault localization is a costly activity in the software development process. Testing and debugging can account for up to 75% of development costs [Tassey, 2002]. In practice, fault localization is performed manually; developers observe failing test cases, and then search the source code for faults. They use their knowledge of the code to investigate excerpts which may be faulty. The

^{*}hamario@ime.usp.br

[†]chaim@usp.br

[‡]fabio.kon@ime.usp.br

most frequent debugging practices include inserting print statements and breakpoints, checking the stack trace, and verifying failing test cases. Since these manual processes can be expensive and ad-hoc [Jones et al., 2007], approaches that automate fault localization are valuable for software development cost reduction.

Several techniques for automating fault localization have been proposed in the last decades [Agrawal et al., 1995, Jones et al., 2002, Wotawa et al., 2002, Zeller, 2002, Renieris and Reiss, 2003]. These techniques use testing information to suggest which program entities, such as statements, predicates, definition-use associations, and call functions [Renieris and Reiss, 2003], are more likely to be faulty. Using fault localization results, developers can inspect the code to search for bugs.

1.1 Motivation and scope

This survey’s main scope is to analyze fault localization techniques that use dynamic information from test execution, known as Spectrum-based Fault Localization (SFL) or Coverage-based Fault Localization. These techniques have achieved significant results when compared to other fault localization techniques and have a lower overhead.

Despite the growing number of available SFL techniques, they are mostly unknown to practitioners. Many factors explain this. In general, SFL techniques have been evaluated using a set of known programs. In most cases, they are small and contain a single fault per version. In practice, though, developers tackle large programs with an unknown number of faults. To complicate things, these bugs may interact and produce different failures depending on the failing test case. User studies in which developers debug real faulty programs with the support of SFL techniques could shed light on such techniques’ effectiveness and efficiency, but they are scarce. As a consequence, the existing SFL techniques are rarely used to automate software companies’ debugging processes [Parnin and Orso, 2011].

This survey presents a comprehensive view of state-of-the-art SFL techniques proposed from 2005 to February 2016, describing the most recent advances and challenges, which includes: approaches and testing information used by SFL techniques; the number and characteristics of faults; benchmark programs used in experiments; costs of SFL techniques; new ways to provide fault localization results to developers; and practical use of SFL techniques.

We discuss and summarize features, limitations, and challenges, indicating future directions for improving SFL. The techniques are classified according to their debugging support strategies and relevance to the main issues. We also present a concept map [Novak and Cañas, 2008] regarding the SFL area, addressing the relationships among topics, their roles, and challenges to future research. We believe the information in this survey is useful to both researchers interested in understanding and improving fault localization techniques, especially Spectrum-based Fault Localization, and practitioners interested in improving their debugging processes.

This paper’s remainder is organized as follows. Section 2 presents an overview of the fault localization area, including history, terminology, and the issues addressed in the survey. We describe the scope, criteria for paper selection, and the issues regarding fault localization in Section 3. In Section 4, we present different fault localization approaches, focusing on spectrum-based techniques. Issues regarding faults are shown in Section 5. The benchmarks used to evaluate the techniques are presented in Section 6. Section 7 shows the issues related to the use of testing information in fault localization. The practical use of fault localization is presented in Section 8. We discuss main challenges and future directions in Section 9. Related works are shown in Section 10. Finally, we present our conclusions in Section 11.

2 Concepts and seminal studies

In this section, we define the main terms used by fault localization studies and present a historical overview of seminal works.

2.1 Terminology

Due to the growth of fault localization studies, several terms have been used to define similar concepts. In what follows, we clarify terms and synonyms used across the studies addressed in the survey.

Faults, errors, and failures

Faults, errors, and failures represent three stages in a program’s execution during which an unexpected behavior occurs. The IEEE Standard 610.12 [IEEE, 1990] defines fault, error, and failure as follows. *Fault* is an incorrect step, process, or data definition in a computer program. A fault is inserted by a developer who wrote the program. A fault is also called *bug* or *defect*.

Error is a tricky term, which is also sometimes used to refer to as a fault, failure, or mistake. In its particular sense, an error is the difference between a computed value and the correct value [IEEE, 1990]. The term is often used to indicate an incorrect state during the program’s execution. Thus, an error occurs when an executed fault changes the program state. Other terms used to express error are *infection* and *anomaly*.

Failure describes a system’s inability to perform its function at the expected requirements [IEEE, 1990]. A failure is observed as an unexpected output, which occurs when an error in a program state leads to a wrong output. A synonym for failure is *unexpected behavior*. *Crashes* are failures that interrupt program executions and thus have an apparent behavior. *Mistakes* are human actions that produce faults [IEEE, 1990].

Real and seeded faults

The literature on fault localization refers to two categories of faults. *Seeded faults* are those intentionally inserted for monitoring detection [IEEE, 1990]. Faults can be manually inserted for experimental purposes, or by using mutation testing. Fault seeding is also known as *fault injection*. Conversely, *real faults* are those that naturally occur during software development.

Ranking metrics

Ranking metrics are used in fault localization to calculate the likelihood that program entities will be faulty. The studies on fault localization use different terms to refer to ranking metrics: *technique*, *risk evaluation formula*, *metric*, *heuristic*, *ranking heuristic*, *coefficient*, and *similarity coefficient*.

Program entity

A *program entity* is a part of a program that can be named or denoted. It comprises any granularity of a program, from a statement to a subprogram. Program entities comprise *statements*, *blocks*, *branches*, *predicates*, *definition-use associations*, *components*, *program elements*, and *program units*.

Spectrum-based fault localization

Program spectra [Reps et al., 1997], also known as code coverage, can be defined as a set of components covered during test execution [Renieris and Reiss, 2003]. *Spectrum-based fault localization* uses information from program entities executed by test cases to indicate entities more likely to be faulty. There are several synonyms of program spectrum used in the literature, such as *code coverage*, *testing information*, *dynamic information*, *execution trace*, *execution path*, *path profile*, and *execution profile*.

Suspicious program entities

Program entities more likely to contain faults are called *suspicious*, *suspected*, *candidate*, and *faulty elements*.

2.2 A Brief History of Debugging Techniques

Unfortunately, developing programs without making mistakes is nearly impossible. Therefore, debugging is an inherent programming activity. The use of the word *bug* originates in Thomas Edison’s time. It was used to indicate flaws in engineering systems [Kidwell, 1998]. In the late 1940s, the *Mark II* computer at *Harvard University* suddenly stopped. Technicians found that a dead moth had shorted out some of the computer’s circuits, and taped the bug into the machine’s logbook [Kidwell, 1998]. The term *debug* was then associated with the activities of finding and fixing program faults. According to Araki et al. [1991], the most primitive debugging practice entails inserting print statements in the code to verify the state of variables.

Despite advances, in practice fault localization has changed little over time. Most of the techniques used by developers today were proposed in the 1960s [Agrawal and Spafford, 1989], and earlier debugging tools originate from the late 1950s [Evans and Darley, 1966]. Some examples are *Gilmore’s debugging tool* (called *TX-O Direct Input Utility System*) [Gilmore, 1957], *FLex Interrogation Tape (FLIT)* [Stockham and Dennis, 1960], and *DEC (Digital Equipment Corporation) Debugging Tape (DDT)* [Kotok, 1961]. The TX-O Direct Input Utility System influenced subsequent, more sophisticated debugging tools. It was based on the idea of moving the debugging program to the computer’s memory, making it possible to verify and modify registers during the execution. As the first tool to implement the concept of breakpoint, FLIT also allows modifying the program during its execution. DDT evolved from FLIT for the *PDP-1* computer. Another advance in the 1960s debugging tools was the conditional breakpoint, which permits a breakpoint’s execution only when it reaches some specific condition. The first tools to provide code tracing were those for high-level languages, such as debugging tools for Lisp and Prolog. Beyond code tracing, debugging tools for high-level languages do not present any additional features compared to those for assembly [Evans and Darley, 1966]. Indeed, the debugging tools used in industrial settings today do not differ much from the above described techniques.

Nevertheless, several techniques have been proposed for automating debugging, most of them for fault localization. Balzer [1969] presented a tool called *EXtendable Debugging and Monitoring System (EXDAMS)*. EXDAMS was one of the first tools to allow either backward or forward navigation through the code. The visualization provides control-flow and data-flow information using graphics, such as a tree structure of the execution at some point of interest. The execution data is stored in a history tape. However, EXDAMS does not use this information to suggest statements more likely to contain bugs, which would help developers in fault localization. Nagy and Pennebaker [1974] proposed one of the earliest techniques to automatically identify bugs by comparing successive

versions of a program, considering that the changes in code are bug corrections. Bug patterns are described as the study’s result.

Some previous techniques tried to understand a program’s whole behavior. These techniques depend on program’s correct specifications, which in practice are very difficult to obtain. Adam and Laurent [1980] proposed a tool called *LAURA*. The tool receives a program model represented by graphs. To identify faults, LAURA makes program transformations to compare the original program with the program model. Johnson and Soloway [1985] proposed a tool called *PROgram UnderSTanding (PROUST)*. PROUST receives *programming plans*, the intentions that a developer has to write the code, and a program. PROUST has a programming plans knowledge base that it compares to the input plan’s goals. PROUST then generates a diagnostic output of bugs found in the code, including an explanation of mistakes causing the bugs.

Assertions are another strategy used to automate debugging. Fairley [1979] proposed a debugging tool called *Assembly Language Assertion Drive Debugging INTERpreter (ALADDIN)*, which uses breakpoint assertions instead of breakpoint locations. The breakpoint assertion is executed when a wrong value occurs in the program state at a certain point. Assertions can be helpful for detecting errors in some circumstances, especially for functions that calculate values or must contain a certain number of elements. However, it is not always possible to use assertions to detect all incorrect program behaviors, which may be infeasible for large and complex programs.

Shapiro [1983] proposed two algorithms to detect incorrect program procedures: one for deterministic programs, and another for non-deterministic programs. These algorithms ask an oracle if the output for a given input to a procedure is correct, repeating the process for the following program execution procedures. The first procedure with an incorrect output is the incorrect procedure. The algorithms proposed by Shapiro [1983] suppose that developers perform the role of the oracle for each executed procedure, which may be an error-prone and time-consuming activity for long-running and large programs. The author suggests that a possible way to automate the oracle is to accumulate a knowledge database of developers’ answers. Fritzson et al. [1992] implements such an idea using category partition testing.

Korel [1988] proposed a tool called *Program Error-Locating Assistant System (PELAS)*—the first fault localization technique based on program dependence. PELAS asks developers about a behavior’s correctness, and uses the answer to indicate possible fault locations. Program dependence data is used to guide the developers’ navigation in searching possible fault sites. The author states that the backtracking reasoning used (based on program dependence) is an abstraction of experienced developers’ intuitive processes. PELAS can narrow the amount of code a developer must verify.

Program slicing, a technique based on programs’ static information, was proposed by Weiser [1981]. The technique generates subsets of a program, called slices, which contain expected program behavior. Thus, a developer focuses his/her attention on a reduced part of a program. Program slicing can be used for debugging, or to change code, depending on the specification of program elements or variables of interest, called *slicing criterion*. A slice is composed of elements that relate to such a criterion. Korel and Laski [1988] proposed *dynamic slicing* to reduce slice size. Dynamic slices are composed of only executed statements. Although dynamic slicing reduces the amount of code to be inspected, the remaining code is still excessive, which makes it impractical.

The use of testing information for fault localization has grown over the last few decades. Collofello and Cousins [1987] proposed the first fault localization technique that uses paths executed by tests to indicate faulty sites. They used ten ranking metrics to calculate the likelihood that program elements will be faulty, wherein a program element is a decision-to-decision path (DD-path)—the code chunk between two predicates. The technique uses a set of DD-paths from passing test cases, and DD-paths from a single failing test case, to indicate DD-paths likely to contain faults.

Agrawal et al. [1995] proposed *execution dices* for fault localization. An execution dice is a set of program basic blocks formed by the difference between two dynamic slices—one from a failing test case and the other one from a passing test case. Even reducing the amount of code returned, the dices still contain a large amount of blocks for inspection.

Other approaches were proposed in the early 2000s. Jones et al. [2002] present a technique called *Tarantula* that uses a ranking metric (also called Tarantula) to calculate statements’ suspiciousness. The suspiciousness values are calculated according to the frequency of the statements in passing and failing test cases. These statements are classified and shown in a graphic visualization form, using different colors according to their suspiciousness values. Zeller [2002] applied the *Delta Debugging* (DD) algorithm to find causes of failures that occur during execution, using the difference in program states (variables and values) between one passing and one failing run. These differences are reduced to obtain a minimal set that causes the failure. Such a subset is deemed the failure’s cause. DD differs from other works by using program states instead of program elements.

Some techniques proposed for fault localization use information from static analysis. These techniques are independent from testing, and can be used to inspect all paths in a program. Static analysis advantageously assures that a program is fault-free by exploring all its possible interactions. However, the performance of these techniques is generally tied to formal proofs of program correctness. Such proofs are infeasible for practice, even for small general purpose programs. Wotawa et al. [2002] used *Model-Based Diagnosis* (MBD), which was originally used for electronic digital circuits, for debugging software faults. MBD generates a logical model from the source code, and uses logical reasoning to obtain a minimal set of statements that explain¹ the existing faults. Hovemeyer and Pugh [2004] proposed a tool that uses bug patterns from Java to automatically locate bugs. The tool statically analyzes a program to search for bug patterns; they proposed fifty bug patterns. The technique generates a list of warnings (statements that might be faulty). Rutar et al. [2004] discuss tools that use static analysis to automatically locate bugs.

In this brief historical overview, we described many different approaches to improve the localization of program faults. However, print statements and symbolic debuggers are prevalent in practice. What is the reason for such a state of affairs in debugging? Primarily, many of the techniques do not scale to programs developed in industry. Another reason is that the techniques are not assessed *in situ*; that is, in real debugging situations. Parnin and Orso [2011] showed that the developer’s behavior may differ from that expected by fault localization technique’s proponents.

The rest of this survey is dedicated to Spectrum-based Fault localization (SFL). SFL utilizes testing information to highlight suspicious pieces of code. Because it utilizes data already collected during testing, SFL tends to have lower overhead in comparison to other debugging techniques. We discuss the more relevant results and challenges to industry adoption of SFL in the following sections.

3 Selection of studies and scope

This survey presents a comprehensive overview focused on techniques that use SFL. Studies published in the last decade, from 2005 until February 2016, were searched in the following digital libraries: *ACM Digital Library*, *IEEE Xplore Digital Library*, *SpringerLink*, and *SciVerse Scopus - Elsevier*.

The studies included in this survey were published in journals and conferences with acknowledged quality. We combined database searches with backward snowballing [Jalali and Wohlin, 2012] to expand the search for relevant results. We selected studies that proposed new techniques to perform

¹*Explain* means components that are logically related to faulty behaviors




Figure 1: Fault localization issues

fault localization based on program spectrum data. Only works that carried out an experimental evaluation of the proposed technique were included. We also considered studies that compare existing fault localization techniques, that propose improvements to program spectra data, and that assess practical use of SFL techniques. When available, only the extended versions of the papers were analyzed.

We selected papers after reading the title and abstract of all studies returned. In doubtful cases, other sections of the papers were read to decide whether they should be included. Selected papers were read in full. For the backward snowballing process, we first selected papers based on the description of such works from source papers. Then, we applied the same criteria to select the most relevant studies. In total, 122 papers were included in this survey. Other relevant works would be included. However, due to space restrictions we do not describe all selected papers.

Despite all efforts and attempts to improve fault localization, developing industrial-level techniques is still a significant challenge. Figure 1 shows the issues that represent the main challenges addressed in this survey. The issues were classified in five major topics: *Techniques*, *Faults*, *Benchmarks*, *Testing information*, and *Practical use*. Most of the studies intersect more than one topic. For each topic, we present studies that present the most distinguishable contributions. Sections 4 to 8 present the studies according to the issues presented above.

4 Fault localization techniques

Spectrum-based Fault Localization techniques propose several strategies to pinpoint faulty program entities. Most of them rank suspicious entities by using ranking metrics, which are based on similarity coefficients and statistical techniques. Artificial intelligence approaches are also used for fault localization. Other SFL techniques are based on execution models that indicate suspicious entities by comparing passing and failing executions. Some studies combine different techniques, while others compare their effectiveness.

There are SFL techniques that make use of other program analysis information, like program dependencies, execution graphs, and clustering of program entities. SFL techniques use different levels of program entities to pinpoint possible faulty sites, from basic blocks to method/function

Table 1: Ranking metrics for fault localization

Ranking metric	Formula	Ranking metric	Formula
Tarantula	$\frac{\frac{c_{ef}}{c_{ef}+c_{nf}}}{\frac{c_{ef}}{c_{ef}+c_{nf}} + \frac{c_{ep}}{c_{ep}+c_{np}}}$	Ochiai	$\frac{c_{ef}}{\sqrt{(c_{ef}+c_{nf})(c_{ef}+c_{ep})}}$
Jaccard	$\frac{c_{ef}}{c_{ef}+c_{nf}+c_{ep}}$	Zoltar	$\frac{c_{ef}}{c_{ef}+c_{nf}+c_{ep}+10000 \cdot \frac{c_{nf}c_{ep}}{c_{ef}}}$
O^p	$c_{ef} - \frac{c_{ep}}{c_{ep}+c_{np}+1}$	Minus	$\frac{\frac{c_{ef}}{c_{ef}+c_{nf}}}{\frac{c_{ef}}{c_{ef}+c_{nf}} + \frac{c_{ep}}{c_{ep}+c_{np}}} - \frac{1 - \frac{c_{ef}}{c_{ef}+c_{nf}}}{1 - \frac{c_{ef}}{c_{ef}+c_{nf}} + 1 - \frac{c_{ep}}{c_{ep}+c_{np}}}$
McCon	$\frac{c_{ef}^2 - c_{nf}c_{ep}}{(c_{ef}+c_{nf})(c_{ef}+c_{ep})}$	Kulczynski2	$\frac{1}{2} \left(\frac{c_{ef}}{c_{ef}+c_{nf}} + \frac{c_{ef}}{c_{ef}+c_{ep}} \right)$
Filter	$\begin{cases} S_T & \text{if } c_{nf} \neq 0 \\ 0 & \text{otherwise} \end{cases}, \text{ where } S_T = \text{Tarantula.}$		

calls. Other studies propose new spectra to be used for fault localization. In this section, we present SFL techniques, exploring the aforementioned strategies to improve fault localization.

4.1 Similarity-based techniques

Similarity-based techniques use ranking metric formulas to pinpoint faulty program entities. To determine correlations between program entities and test case results, these ranking metrics use program spectrum information derived from testing as input. Each program entity receives a suspiciousness score that indicates how likely it is to be faulty. The rationale is that program entities frequently executed in failing test cases are more suspicious. Thus, the frequency in which entities are executed in failing and passing test cases is analyzed to calculate its suspiciousness score. There are ranking metrics specifically created for fault localization, and other metrics were adapted from areas such as molecular biology. Some studies perform experiments comparing ranking metrics. Works that propose or use ranking metrics are presented hereafter.

Tarantula [Jones et al., 2002] was one of the first techniques proposed for SFL. Its formula is shown in Table 1. For each statement, Tarantula calculates the frequency in which a statement is executed in all failing test cases, divided by the frequency in which this statement is executed in all failing and passing test cases. Table 1 shows some of the main ranking metrics used for SFL. We use the following nomenclature in Table 1: c_{ef} indicates the number of times a statement (c) is executed (e) in failing test cases (f), c_{nf} is the number of times a statement is not executed (n) in failing test cases, c_{ep} is the number of times a statement is executed by passing test cases (p) and c_{np} represents the number of times a statement is not executed by passing test cases.

Tarantula has been used in several studies. Jones et al. [2007] use Tarantula to calculate the suspiciousness score of statements for their parallel debugging technique (see Subsection 5.1). Ali et al. [2009] use Tarantula to evaluate characteristics that can influence fault localization techniques. Furthermore, Tarantula has been used as benchmark by several other techniques [Wong et al., 2010, Xie et al., 2013].

Some studies propose ranking metrics similar to Tarantula for techniques that use other coverage types. Masri [2010] uses a *Tarantula-like* ranking metric for the DIFA coverage (see Subsection 4.8). This ranking metric is combined with another that considers only the percentage of executions in failing test cases. The final suspiciousness score averages the two. Wang et al. [2009] use a ranking metric similar to Tarantula for basic blocks. Chung et al. [2008] propose a ranking metric for predicates that is also similar to Tarantula.

In addition to Tarantula, other techniques based on similarity formulas have been proposed

in the last years. Abreu et al. [2007] propose using *Ochiai* and *Jaccard* similarity coefficients as fault localization ranking metrics. *Ochiai* was originally used in molecular biology, and *Jaccard* was used by Chen et al. [2002] to indicate faulty components in distributed applications. The results of Abreu et al. [2007] show that both *Ochiai* and *Jaccard* outperform Tarantula on fault localization effectiveness. From these results, several works have used *Ochiai* [Santelices et al., 2009, DiGiuseppe and Jones, 2015]. *Jaccard* was not used on its own by any study presented in this survey. However, it was used in studies that compare the performance of several ranking metrics [Naish et al., 2010, Xie et al., 2013, Ma et al., 2014].

Xu et al. [2013] propose a Tarantula-like ranking metric called *Minus KBC* (MKBC) for their KBC coverage (see Subsection 4.8). The difference is that MKBC subtracts the percentage that a KBC is not executed in failing test cases, and divides by the percentage that such KBC is not executed for all executions (see Table 1). This complementary frequency is called *noise*; it decreases the importance of non-execution in the analysis.

Naish et al. [2011] propose two new ranking metrics optimized for single-fault programs, called O and O^p . Assuming the existence of a single fault, only statements that are executed in all failing test cases are fault candidates. *Kulczynski2* (from Artificial Intelligence) and *McCon* (from studies of plankton communities), are ranking metrics that have presented better results for programs with two simultaneous faults [Naish et al., 2009]. Wong et al. [2012b] presented a technique called *DStar*, which is a modified version of the Kulczynski ranking metric. The idea behind *DStar* is that the execution trace of each statement through test cases can be viewed as an execution pattern. Thus, the similarity between statements more frequently executed by failing test cases can pinpoint the faulty ones.

Other techniques propose different strategies in conjunction with ranking metrics. Jeffrey et al. [2008] presented a technique that replaces values of statements in failing executions. If a failing execution then passes, such statement is classified between the most suspicious. The technique uses values observed in a statement from all executions, and only one value for each statement is replaced per execution. Naish et al. [2009] presented an approach that assigns different weights to statements in failing test cases according to the number of statements in an execution. The lower the number of commands, the greater the chance that one of them will be faulty.

Xie et al. [2010] proposed a technique that forms two groups of statements: one with statements that the failing test cases executed at least once, and another with statements that these test cases did not executed. The first group’s statements receive suspiciousness scores. Bandyopadhyay and Ghosh [2011] proposed a technique that assigns different weights for test cases according to the similarity of a passing and a failing test case. The more similar a passing test case is, the higher is its weight.

4.2 Statistics-based techniques

Statistical techniques have also been applied in fault localization, and are used in the same manner as similarity-based techniques. However, each statistical technique uniquely deals with testing information. Liblit et al. [2005] proposed using conditional probability to evaluate predicates during program executions for fault localization. Predicates that are evaluated only as true in failing executions are deemed as more suspicious. They calculate the probability that a predicate p with value true causes a failure ($Failure(p)$), and the probability that an execution of p causes the failure ($Context(p)$). The difference between such values (*Importance*) indicates the suspiciousness of each predicate. Liu et al. [2006] proposed a statistical fault localization technique called SOBER that considers the results of predicates evaluated as true multiple times per execution for several executions. They used *Bernoulli distribution* to model each predicate’s result for each execution. Then, conditional probability for passing and failing executions is used to calculate a predicate’s

bug relevance.

Nainar et al. [2007] proposed the idea of complex predicates as bug predictors for fault localization. Complex predicates are formed by two predicates that are evaluated in each execution using boolean function operations (conjunction and disjunction). They argued that combinations of predicates can enhance fault localization when the predicates that form the complex predicates are already good bug predictors. Baah et al. [2010] studied the use of causal inference for fault localization, aiming to enhance fault localization by isolating causal effects that occur between program entities in a program dependence graph. This isolation can improve the values assigned to faulty entities by reducing noise caused by other program elements in the presence of failures. They used causal graphs concepts to identify entities that are independent in a program dependence graph. They also proposed a linear regression model to calculate the causal effect of statements on failures.

Zhang et al. [2011] proposed a technique for using a non-parametric statistical model for fault localization. They observed that the distribution of predicates during executions is non-normal (see Subsection 4.8). To calculate the suspiciousness of predicates, they consider the difference between the probability density function of a random variable of passing test cases and a failing test case. Wong et al. [2012c] use a crosstab-based technique for fault localization. The chi-square test was used to calculate the null hypothesis that an execution is independent of the coverage of a statement. Chi-square is applied to deal with categorical variables; in this work such values are the test case results (pass and fail), and the presence of a statement in an execution (executed or not). They measured the dependency between an execution and a statement.

Zhang et al. [2012] used maximum likelihood estimation and linear regression to tune in SFL lists. They used previously known lists and bug positions, assuming a symmetric distribution of bug positions in such lists. Thus, they estimated a position to adjust the lists. Xue and Namin [2013a] applied the *odds ratio* to SFL. The odds ratio is a statistical technique often used in classification and text mining problems. The odds ratio measures the strength or the weakness of a variable (a statement executed or not executed) associated with an event (a passing or a failing test case).

4.3 Artificial intelligence-based techniques

Artificial Intelligence (AI) techniques have been applied for fault localization, using program spectrum data as input for classifying suspicious program elements.

Liu et al. [2005] proposed a technique that uses *graph mining* and *support vector machines* (SVM) for fault localization. Program executions are represented as behavior graphs. Each node of a behavior graph is an executed function. The graphs are labeled with their execution result. Graph mining is applied to discover frequent subgraphs, which are features used to classify the graphs. SVM is applied to classify incorrect and correct executions. Thus, a sequence of bug-relevant functions is presented for a developer’s inspection. Nessa et al. [2008] applied *N-gram analysis*, from the data mining, for fault localization. The technique generates N-grams, subsequences of statements, from program spectra. The technique uses *Association Rule Mining* to calculate the conditional probability that each N-gram relates to faulty executions. A list of the most suspicious statements is obtained from the most suspicious N-grams.

The technique proposed by Murtaza et al. [2008] uses a *decision tree* as a heuristic to indicate the origin of a fault. The technique identifies patterns of function calls related to the fault. Wong et al. [2012a] developed a *Radial Basis Function neural network* (RBF-NN) for fault localization. The RBF-NN is modeled as statements representing the input neurons. Thus, the neural network acts as a ranking metric, in which the neural network individually evaluates each statement to determine its suspiciousness. Lucia et al. [2014] assessed the use of association measures for fault localization. They evaluated 20 association measures commonly used in data mining and statistics, such as *Yule-*

Q , $Yule-Y$, and *odds ratio*. They modeled the problem as the strength of association between each entity’s execution (and non-execution) and failures. Two association measures had a performance comparable to Ochiai: *information gain* and *cosine*.

Roychowdhury and Khurshid [2011] proposed a technique based on *Feature Selection* from Machine Learning. They used two well-known methods, *RELIEF* and *RELIEF-F*, to classify statements according to their relevant potential to be faulty. The coverage matrix obtained from the test execution is applied to RELIEF and RELIEF-F. Each executed statement is considered a feature, and each coverage of a certain test case is a sample. This techniques aim is to capture the diversity of these statements’ behaviors as they relate to bug execution. Zhang and Zhang [2014] applied *Markov Logic Network* from machine learning to fault localization. They modeled the problem by considering dynamic information (program spectra), static information (control-flow and data-flow dependence), and prior bug knowledge (locations of similar bugs found in the past). The technique is developed for single-bug programs.

4.4 Combination of similarity-based techniques

The many techniques based on ranking metrics led to studies that try to combine them to improve fault localization. Debroy and Wong [2011b] used Tarantula, Ochiai, and Wong3 to propose a consensus technique, using the concept of rank aggregation, specifically *Borda’s method*, to combine statements classified with different values by each previous technique. Ju et al. [2013] proposed a new ranking metric called *HSS*. They combined two ranking metrics, O^p and Russel&Rao, by multiplying them to calculate the suspiciousness of program entities. Xie et al. [2013] show both O^p and Russel&Rao to be better-performing ranking metrics (see Subsection 4.5).

Other studies use AI to create or combine previous techniques. Yoo [2012] built several ranking metrics using *Genetic Programming* (GP). The GP operators used to create such ranking metrics were simple arithmetic operations: addition, subtraction, multiplication, division, and square root operation. To measure how well a ranking metric classifies faults, they used the EXAM score evaluation metric (see Subsection 8.1) as the fitness function. Six out of 30 ranking metrics created using GP outperformed other ranking metrics used for comparison. Cai and Xu [2012] proposed a technique that uses suspiciousness lists from previous ranking metrics. The lists of 28 ranking metrics were used as input. The technique uses a *k-means* algorithm to cluster the statements, using the ranking position obtained by each ranking metric as features of the statements.

The technique proposed by Le et al. [2015b] combines SFL (Tarantula) and Information Retrieval-based fault localization (IRFL). IRFL uses bug reports to generate suspiciousness lists. Their technique combines program elements from failing test cases with related suspicious words. The results are present at method level.

4.5 Comparison of similarity-based techniques

While some studies use several similarity-based techniques to evaluate their proposed techniques, others compare techniques to identify which are more effective. Naish et al. [2009] use 11 ranking metrics to perform experiments with their technique. The technique assigns different weights for statements according to the number of statements executed in failing test cases; among them are Tarantula, Ochiai, Jaccard, and Wong3. The ranking metrics O^p , Wong3, and *Zoltar* were more efficient for programs with single faults. Debroy and Wong [2011a] show the equivalence between different ranking metrics by comparing and simplifying their formulas. These ranking metrics classify program entities in the same relative position in their ranking lists.

Xie et al. [2013] performed a theoretical comparison between 30 ranking metrics used in fault localization. To make this comparison, they grouped statements classified by the ranking metrics according to statements with respective higher, equal, and lower scores than the faulty statement. They identified six groups of equivalent ranking metrics, and seven ranking metrics that lack equivalent metrics. They also showed that some ranking metrics perform hierarchically better than others. Le et al. [2013] evaluated the study of Xie et al. [2013] using well-studied benchmarks, showing that the theoretical comparison does not hold for such programs.

Ma et al. [2014] compared seven ranking metrics for SFL. To model the comparison, they proposed a *Vector Table Model* that represents the possible state of each statement for any program. They assumed that a program has a single fault. The results show that O , O^p , $D^E(J)$, $D^E(C)$ are equivalent. These metrics outperform *Wong1*, and also outperform Jaccard and *Kulczynski1* (which are equivalent). $D^E(J)$ and $D^E(C)$ were proposed by Naish and Lee [2013]. Kim and Lee [2014] carried out another ranking metric comparison, classifying 32 ranking metrics using clustering. Similarity was measured using the normalized suspiciousness values of such ranking metrics to compare their effectiveness, and the metrics were clustered in three groups of equivalent heuristics. They pointed out that these groups have complementary characteristics, each of which has its weaknesses and strengths.

4.6 Execution models

Another strategy used by SFL techniques is to build execution models from test executions. Such models represent patterns of failing and/or passing executions. Execution models are used to identify entities that meet or flee from an expected pattern. There are also models represented by graphs of execution. Wang and Roychoudhury [2005] proposed a technique that generates a passing run from a failing run. The technique consists of toggling the results of branches in the failing run until obtaining a passing run. The outcome is a list of branches that were modified in the passing run. The toggling process starts on the last branch executed, and more branches are incrementally toggled until a passing run is located. A developer must check whether the generated run is correct. Zhang et al. [2006] also proposed a technique that toggles branches of a failing run to generate a passing one. However, only one predicate is toggled per execution. The process starts on the last executed predicate, and can be applied backwards in all the following predicates until generating a passing run.

The technique proposed by Zhang et al. [2009b], called *Capture Propagation*, performs the propagation of suspicious values between related code blocks using graph models. The technique generates two mean edge profiles, one for all passing test cases, and another for all failing test cases. These profiles are used to obtain the suspiciousness for each edge. A propagation ratio of an edge is calculated according to the number of edges that enter in a successor block. Finally, the suspiciousness value of the predecessor block is obtained from the sum of the propagation ratio of all successor blocks for which the predecessor block has an edge. The technique presented by Mariani et al. [2011], detailed in the Subsection 4.8, uses a behavior model of method calls from passing test cases. This model is compared with interactions from failing test cases to indicate suspicious interactions. Liu et al. [2010] proposed a technique in which the model is composed of messages from objects within the same class.

Wan et al. [2012] proposed a behavioral model that is constructed using two different coverage types: objects of an OO program and calls occurred inside each object. A model is a sequence of objects and calls in execution traces of failing and passing test cases. The model contains two levels: the objects, and its internal calls. Two values from each entity are used to compose a suspicious value: coverage and violation. The violation means that entities from failing executions that are

not present in the model are more suspicious, and entities which are present in the model from passing executions are less suspicious. Dandan et al. [2014] proposed a probabilistic model of state dependency for fault localization. State dependency relates to predicate statement, which can be true or false. The technique calculates the probability that a predicate is true or false in passing and failing executions. Two probability models are generated: one for passing executions, and another for failing executions. The probability of control dependent statements is based on the probability of their parent statements. The suspiciousness value for each statement is then calculated using the models' probability values, and the outcome is a list of the most suspicious statements.

4.7 Program analysis-based techniques

In their attempts to highlight faulty code excerpts, some techniques use additional program analysis information to highlight faulty code excerpts, such as program dependence data, program slicing, and assignment of weights to differentiate program entities. The technique proposed by Zhang et al. [2009b] (see Subsection 4.6), considers the influence between nodes and related edges from a control-flow graph to assign weights for code blocks. Zhao et al. [2010] propose a technique for calculating the suspiciousness of edges in a control-flow graph. This value is used to obtain a weighted coverage for each edge that considers the distribution of the control-flow for passing and failing test cases. The suspiciousness of each block is then calculated using the passing and failing weighted coverages for each block through a *Tarantula-like* metric.

Yu et al. [2011] propose a technique that uses an *Ochiai-like* metric to calculate the similarity of control and data dynamic dependences. This similarity indicates the correlation between these dependences and an incorrect program behavior. The final suspiciousness scores are obtained by considering the maximum value between the suspiciousness of control and data dependences for each statement. Zhao et al. [2011a] proposed a technique that uses the influence of edges (branches) in an execution to obtain a list of suspicious blocks. First, they calculate the suspiciousness of edges. Afterwards, they calculate the *fault proneness* of each edge using the total nodes that *in* in the successor node and total nodes that *out* from the predecessor node. The approach indirectly uses the dependencies between basic blocks to measure the importance of predecessor blocks in their successor blocks.

For their work on fault localization for field failures, Jin and Orso [2013] proposed three strategies to reduce the amount of entities (branches) presented for fault localization: (1) *filtering*, which excludes entities at three levels: those executed only by passing executions, branches executed for both passing and failing executions, and branches that have other control dependent branches; (2) *profiling*, which computes the suspiciousness of program entities by considering the number of times an entity is executed in each execution for all executions; and (3) *grouping*, which groups entities that belong to the same code region and have the same suspiciousness values. These strategies can in most cases pinpoint the faulty entities in the first picks.

Li et al. [2014] proposed a technique that uses information from program structure to improve fault localization. They consider that some structures can influence their statements, making them more susceptible to be wrongly classified by SFL techniques. For example, a faulty statement in a main class can be underestimated because it is always executed by both failing and passing runs. Conversely, a non-faulty statement in a *catch* block can be overestimated because it is only executed by failing runs.

Some studies use program slicing because of their ability to hold faulty statements in their resultant subsets. However, slicing-based techniques must propose strategies for reducing the amount of returned code. Zhang et al. [2005] evaluated three variations of dynamic slicing for fault localization: data slicing, full slicing, and relevant slicing. Regarding the amount of information returned,

relevant slices are larger than full slices, which are in turn larger than data slices. The results show that, although data slices returned a reduced amount of statements, in several cases the fault statement was not returned. Full slicing returned a large amount of statements, with the faulty statement in most cases among them. Relevant slicing returned a slightly larger amount of statements than full slicing, but the faulty statement was always present. Wong and Qi [2006] proposed a technique that uses execution slices and inter-block data dependency for fault localization. The technique first computes a dice from one failing and one passing execution. If the fault is not in the dice, inter-block data dependency from the failing test case is used to add more information. If the amount of code after adding such data dependencies is excessive, the technique uses the distinct blocks presented in another passing execution, with respect to the previous passing execution, to generate another dice.

Alves et al. [2011] used dynamic slicing and *change impact analysis* to reduce the number of statements in the ranking lists provided by fault localization techniques. Three techniques were proposed to obtain this reduction: *T1*, which ranks only statements that appear in the dynamic slice; *T2*, which applies a change impact analysis before the dynamic slicing and considers all statements in the dynamic slice; *T3*, which is similar to *T2*, but with only changed statements ranked. Tiantian et al. [2011] proposed a technique that calculates suspiciousness of predicates using SFL. F-score is used to assign values that indicate the likelihood that predicates will be faulty. From a predicate, the technique generates control-flow and data-flow information on demand for each predicate by constructing a procedure dependence graph (PDG) for the procedure that includes the predicate. Next, backward and forward slices from this PDG are obtained.

Ju et al. [2013] proposed a technique that combines full slices from failing executions with execution slices from passing executions. These slices are merged using intersection to obtain a hybrid slice. Wen et al. [2011] proposed a technique that combines SFL and program slicing; it removes all programs elements that do not belong to any failing executions slices. To calculate the suspiciousness of the remaining elements, they consider the execution frequency of each element in each test case, and the contribution of an element in a test case. The contribution is the percentage in which an element is executed, considering all executed elements.

4.8 Program entities

Program spectra can represent different levels of program structures, from fine-grained to coarser excerpts. The different coverage types include statements, blocks, predicates, function or method calls, and others developed by a new technique. Some studies also combine different coverage levels.

The program entities most commonly used are statements. The examination of statements may lead a programmer directly to the faulty site. However, a large number of statements may need to be examined before the faulty statement is found. Several works have used statements as their program entity investigation level [Jones et al., 2007, Xie et al., 2010, Le et al., 2013]. Other works used block coverage [Wong and Qi, 2006, Zhang et al., 2009b, Xue and Namin, 2013b].

Predicates are statements such as branches, returns, and scalar-pairs [Liblit et al., 2005]. Guo et al. [2006] use predicate coverage to compare one failing execution to all passing executions, and then to identify the most similar passing execution. This similarity is measured by comparing the results of predicates (true or false) and their execution order. The technique generates a report composed of predicates that were executed only by the failing execution. Naish et al. [2010] propose a technique that generates predicate information from statement coverage to perform fault localization. The authors show that predicate coverage provides more information about the execution, such as execution results and control-flow data, than statement coverage. Zhang et al. [2009a] perform a statistical study of the distribution behavior of predicates that are relevant to the occurrence

of failures. In their experiments, about 40% of predicates lack normal distribution. Chilimbi et al. [2009] proposed a technique that uses path profiles (intra-procedural code segments) and conditional probability to identify paths that are more likely to be faulty. Their results showed that paths are more precise than predicates for pinpointing faults.

Dallmeier et al. [2005] used method call sequences in their technique. The method call sequences that income and outcome in an object are summarized to represent a sequence set per class. Classes whose sequence sets differ from one failing run to several passing runs are more suspicious of being faulty. In this work, the authors observed that incoming calls are more likely to contain faults than outgoing calls. Laghari et al. [2015] proposed an SFL technique that also produces suspiciousness lists at class level. Mariani et al. [2011] present a technique that gathers information about interactions (sequences of method calls) between software components². The technique generates an interaction model of passing test cases. This model is compared to interactions from failing test cases to indicate suspicious method calls. Other works also used method call coverages [Yilmaz et al., 2008, de Souza and Chaim, 2013].

New coverage types have been proposed for fault localization. Santelices et al. [2009] compared the performance of different coverages—statement, branch, and dua. They showed that different faults are best located by different coverage types. Three approaches were proposed combining such coverages. One of them (*avg-SBD*), which calculates the average suspiciousness values of statement, branch, and dua coverages, achieved better results. In this same study, the authors used a technique that infers an approximation of dua coverage (*dua approx*) using branch coverage data, which demands lower execution costs. Yilmaz et al. [2008] proposed a technique based on the concept of time spectra. The technique collects the execution times of methods in passing runs to build a behavior model for each method. Methods executed in failing runs that deviate from the model are considered more suspected. As time is an aggregate value, it can also represent sequences of events.

Masri [2010] uses a coverage named *Dynamic Information Flow Analysis* (DIFA) for fault localization. DIFA is composed of interactions performed in an execution, including data and control dependences from statements and variables. These interactions are known as *information flow profiles*.

Xu et al. [2013] present the coverage *Key Block Chains* (KBC). Each KBC is composed of only one predicate whose execution result is true, known as an atomic predicate. Thus, a KBC may have several sizes, according to the code blocks that are executed until a predicate is evaluated as true. Papadakis and Le Traon [2013] use mutants as new type of coverage for fault localization. They calculate the suspiciousness of mutated versions using Ochiai. Next, they map the mutants for the statements using the location in which they were inserted.

Strategies to reduce coverage data gathered can also contribute to enhancing fault localization. Perez et al. [2012] propose a technique that reduces the amount of code coverage information by choosing the granularity of the inspected elements that are collected. Only the most suspect elements in a coarser level are collected at a more fine-grained level. The proposed strategy reduces the average execution time for performing fault localization. The technique presented by Wang and Liu [2015] performs multiple predicate switching to find predicates that are critical for revealing faults. To avoid exponential growth due to the number of predicates, they first apply Ochiai to reduce the number of predicates on the switching phase.

²Components in this study of Mariani et al. [2011] are independent programs with specifics functionalities used by other programs

5 Faults

In practice, a developer does not know how many bugs a program has; it may contain single or multiple faults. Simultaneous faults may change the execution behavior, causing interferences in the test results and affecting the performance of SFL techniques. Faults can present different characteristics, also impacting in SFL techniques’ suspiciousness lists. We present these issues regarding faults below.

5.1 Single and Multiple faults

Most Spectrum-based Fault Localization techniques are assessed using programs with single faults. However, a few techniques were proposed specifically for programs with single faults [Abreu et al., 2011, Naish et al., 2011, Zhang and Zhang, 2014], whereas most techniques do not specify such a limitation, yet were assessed using single-fault programs [Guo et al., 2006, Mariani et al., 2011]. For experimental purposes, in which the faults are already known, it is acceptable to suppose that a program has only one fault. However, it is impossible to foresee how many faults a program has in industrial settings. Concerning this issue, SFL techniques started to be evaluated in programs with multiple faults. The following techniques were proposed addressing multiple faults or using multiple faults in their evaluation.

Zheng et al. [2006] used a bi-clustering process to classify predicates and failing runs: they calculated the conditional probability of predicates related to failing runs. The selected predicates were applied in a collective voting process, in which each failing run votes for its favorite predicate. This voting process is iterative, aiming to identify strong relationships between predicates and failing runs.

Another strategy is to identify similar failing test cases. Jones et al. [2007] introduced a technique to parallelize debugging for multiple faults by grouping failing test cases that are most similar to a particular fault. They use *agglomerative hierarchical clustering* to obtain these groups, called *fault focusing clusters*. DiGiuseppe and Jones [2012a] investigated factors that can affect failure clustering. They show that execution profiles and test case outputs can be accurately used for failure clustering. They also verified that failures whose profiles and outputs differ due to the presence of multiple faults can impair the failure clustering process. Such failures may be pre-processed to improve failure clustering. In a following study, DiGiuseppe and Jones [2012b] presented a failure clustering technique that uses semantic information from source code to improve the failure clustering process.

Liu et al. [2008] compared six failure proximity techniques. These techniques are used to classify failure reports that are related to the same bugs. A failure proximity technique extracts a failure’s signature and uses a function to calculate the distance from other failures to group them as related to the same fault. Högerle et al. [2014] investigated factors that impact debugging in parallel for multiple faults. They highlight the infeasibility to obtain pure failure clusters. Such clusters should contain entities and failing test cases entirely related to each fault. One trade-off is obtaining clusters that share program entities, causing what they called *bug race*. Bug race means that faults can be present in more than one cluster. Another trade-off is obtaining clusters that share tests instead of program entities. Bug races are avoided in this case, but some clusters may have no faulty entity.

The work proposed by Dean et al. [2009] uses linear programming to locate multiple faults. The technique returns a set of statements that explain all failing test cases. For the experiments, the single-fault versions of the *Space* and *Siemens suite* were joined to create a version with all faults per program. Naish et al. [2009] conducted experiments with several ranking metrics in programs with one and two faults per version. The results show that some ranking metrics were

more effective for programs with two faults, while others were more effective for single-fault programs (see Subsection 4.1).

Steimann and Bertschler [2009] proposed a technique that uses only failing test cases, assuming that in each of them there is at least one program entity (method) that explains the fault. All entities in these test cases have the same probability of being faulty. All possible combinations between methods through the test cases are verified, and only combinations that can explain the fault are kept. Some entities can be more present than others in the remaining explanations. Thus, the entities are classified according to their frequencies in the explanations.

Abreu et al. [2011] proposed a technique for multiple faults called Zoltar-M. The authors used MBD (Model-based Diagnosis) along with program spectra information to obtain groups of states that relate to the existing faults. The MBD is constructed under logical propositions from the results of the dynamic execution information. Yu et al. [2015] presented a technique for distinguishing test cases that fail due to single faults from those that fail due to multiple faults. The technique creates test sets composed of one failing test case and all passing test cases. Then, they calculated the distance between a failing test case and its most similar passing test case. With the presence of multiple faults, tests with large distances are more likely related.

5.2 Fault interference

The concern about the occurrence of multiple faults led to studies of interferences between simultaneous faults. Debroy and Wong [2009] studied the existing interference between multiple faults. They showed that simultaneous faults can cause interferences in which some faults hide the incorrect behavior of other faults. Conversely, some faults can collaborate to manifest failures related to other faults. The experiments performed showed an incidence of fault interference in 67% of the assessed programs.

Xue and Namin [2013b] studied the impact of multiple faults on five fault localization techniques. They showed that fault interference can reduce fault localization effectiveness by 20% using Ochiai. They also demonstrated improved fault localization effectiveness in around 30% of the cases. Regarding the ranking metrics used, they observed that some of them, like Tarantula and Ochiai, poorly performed as the number of faults increased. The ranking metrics Chi-Square and Odds Ratio exhibited no difference in their performance as the number of faults increased.

DiGiuseppe and Jones [2015] investigated the influence of multiple faults on the performance of SFL techniques. The results showed that the presence of multiple faults has little impact (a decrease around 2% of effectiveness) on the performance of SFL techniques to find the first fault. The fault localization effectiveness of the other faults (beyond the first fault) is also impaired as the number of faults grows. They also showed that the suspiciousness scores of faults tend to decrease as the number of faults increases. There were cases of improved effectiveness, and other cases in which some faults became unlocalizable. They also verified that fault interference occurred in 80% of the assessed programs.

5.3 Fault classification

A few studies have addressed the impact of fault types on their techniques. Santelices et al. [2009] cite that different coverages can contribute to locating distinct fault types, but do not present any relationship between the proposed coverages and the fault types that such coverages locate better. Guo et al. [2006] assessed their technique in the presence of three types of faults: branch faults, assignment faults, and code omission.

Faults by code omission are generally difficult to locate [Zhang et al., 2009b, Xu et al., 2011,

Xie et al., 2013]. Zhang et al. [2007] proposed a technique to deal with faults caused by code omission. They used the concept of implicit dependence to identify indirect dependencies between the use of a variable and a previous conditional statement.

There are works that show how the proposed techniques deal with specific faults. Masri [2010] described the faults used in their experiments, and analyzed the influence of these faults' characteristics on his technique. Burger and Zeller [2011] also described the types of faults used in experiments, and discussed their impact on the new technique. Zhang et al. [2011] used a fault classification defined by Durães and Madeira [2006] to verify the frequency of these faults in real programs.

Debroy et al. [2010] pointed out another type of fault that relates to single faults spread over more than one statement, also known as single faults in multiple lines. Thung et al. [2012] examined 374 faults from three real systems to understand when faults are localizable. By localizable, they meant faults present in one or several lines of code in a nearby region, which is the general assumption of fault localization techniques. They manually inspected all faults, finding that 30% of such faults occur in a single line, while around 10% of the faults spread to more than 25 lines each. Less than 45% of these faults appear in a single method, and less than 75% take place in a single file. Thus, fault localization techniques must be able to locate faults scattered across different code regions.

6 Benchmarks

Several programs have been used to assess SFL techniques; some are program suites composed of small programs often used as benchmarks, while others are medium and large programs. Most of them are open source programs included in software testing repositories, such as the Software-artifact Infrastructure Repository (SIR) [Do et al., 2005]. SIR contains C and Java programs prepared for experimental use, including seeded and real faults, and scripts to automate the execution of controlled experiments.

Next, we present a description of programs used as benchmarks by the SFL techniques. We also describe their characteristics, such as size and number of faults.

6.1 Description of the main benchmarks

Several programs have frequently been used to carry out fault localization experiments. Among them are *Siemens suite* [Hutchins et al., 1994], *Unix suite*, *Space*, *flex*, *gcc*, *grep*, *gzip*, *make*, and *NanoXML*. Other programs were used only once. Table 2 shows a description of benchmarks often used in the studies, the average number of lines of code (LOC) per version, number of faults for all versions, number of versions, and the average number of test cases (No. TC) per version. The Siemens suite and Unix suite data show the average of their programs. The number of LOC, faults, versions, and test cases may vary throughout the studies.

The Siemens suite comprises seven small programs written in C: *print_tokens*, *print_tokens2*, *replace*, *schedule*, *schedule2*, *tcas*, and *tot_info*. Each program contains one fault per version and a test suite including thousands of test cases. The large test suites were created to achieve a high testing coverage. To simulate realistic faults, the faults were seeded for experimental purposes by ten people [Hutchins et al., 1994]. A large number of studies used Siemens suite in their experiments [Guo et al., 2006, Dean et al., 2009, Dandan et al., 2014]. The Unix suite is a collection of Unix utilities written in C that are small in size and contain several faulty versions [Do et al., 2005]. These programs are *Cal*, *Checked*, *Col*, *Comm*, *Crypt*, *Look*, *Sort*, *Spline*, *Tr*, and *Uniq*. The faults of Unix suite were seeded using mutation-based injection. Several works have used Unix suite to assess their techniques [Wong et al., 2010, Roychowdhury, 2012].

Table 2: Benchmarks most used by the studies

Program	Description	LOC	Faults	Versions	No. TC
Siemens suite	7 programs	483	19	1	3,115
Unix suite	10 programs	261	17	17	401
Space	Satellite antenna controller	6,200	38	1	13,585
flex	Lexical analyzer	10,459	21	6	567
grep	Search for patterns in files	10,068	18	6	470
gzip	Data compressor	5,680	28	6	211
make	Build manager	35,545	19	6	793
NanoXML	XML Parser	7,646	32	5	216
Ant	Build manager	80,500	18	11	871
gcc	C compiler	95,218	5	1	9,495
XML-security	Encrypter	16,500	52	3	94

Other medium and large-sized programs have been used for fault localization. Such programs can provide more realistic scenarios for experiments, due to their sizes and different domain characteristics. The Space program, used in several studies, was developed by the European Space Agency. It contains 38 real faults discovered during the development. The test suite was created by Vokolos and Frankl [1998] and Aristotle Research Group [2007], consisting of 13,585 test cases to guarantee that each branch is exercised by at least 30 test cases [Jones et al., 2007]. *Flex*, *grep*, *gzip*, and *make* are medium-sized Unix utilities also frequently used for experiments [Liu et al., 2008, Wong et al., 2012a]. Ali et al. [2009] describe the process adopted to prepare the program *Concordance* for use in fault localization experiments. They also argue that the hand-seeded faults of the Siemens suite may not be suitable to represent programs.

NanoXML is an XML parser for Java with different versions used as benchmarks. The program has both real and seeded faults. Ant and *XML-security* are other Java programs used in SFL experiments [Mariani et al., 2011, de Souza and Chaim, 2013].

6.2 Size

A program’s size is generally used to refer to it as a large (or real) program, or as small program. There is no precise definition limit to determine a program as small, medium, or large in size. Zhang et al. [2009b] consider *flex*, *grep*, *gzip*, and *sed*—programs that contain between 6.5 and 12.6 KLOC—as real life medium-sized programs. Other authors consider these programs as large in size [Debroy et al., 2010, Abreu et al., 2011]. Space, with about 10 KLOC, is generally considered a large and real program [Naish et al., 2009, Yu et al., 2011]. Generally, we can assume that programs with more than 10 KLOC are large programs. Benchmarks between 2 and 10 KLOC are medium-sized programs, and programs with less than 2 KLOC are considered small programs.

We identified that most authors consider programs ‘real’ if they are applied to professional use, whereas ‘not real’ programs (called *toy* programs) are those used for experimental purposes or to perform small tasks, including operational system utilities [Naish et al., 2009, Zhang et al., 2011]. Real programs in this context also consider the existence of real program faults. We observed that the programs considered as medium and large in size by the studies in this survey can be assumed as real programs, and the small programs as toy programs.

Some experiments with large programs used only a few parts of them. The *Columba* program, an email client system, contains about 1,700 classes. For the Gcc program, Wong et al. [2012a] instrumented one sub-directory (gcc/cp) for their experiments. Mariani et al. [2011] uses NanoXML,

Eclipse, and Tomcat, analyzing interactions for some components of such programs.

6.3 Number of faults

As discussed in Subsection 5.1, an important issue that should be investigated by SFL techniques is the presence of multiple faults. To achieve such a goal, benchmarks used in the experiments must also have multiple faults. Moreover, some of these faults must change the output behavior of other faults to address the study of fault interferences.

The existing benchmarks are generally composed of single faults. Researchers usually change the subject programs to generate multiple-fault versions for their experiments, yet these modifications can add biases to the evaluation and make experiments more difficult to reproduce. Jones et al. [2007] created 100 versions of Space, each of them containing from 1 to 8 faults, randomly combining the single-fault versions. Abreu et al. [2011] also generated multiple-fault versions for gzip, Space, and sed in their experiments. Several other works have used these random strategies to generate multiple faults [Wang et al., 2009, DiGiuseppe and Jones, 2015].

Other studies identified faults to carry out their multiple-fault experiments. Steimann and Bertschler [2009] claim that the number of available multiple-fault benchmarks is quite limited. In their experiments, they show an example with three simultaneous real faults from the program *Apache Commons Codec*. Wong et al. [2012a] identified five existing faults of the Gcc compiler using the *Bugzilla database*. They merged these bugs to create a 5-bug version of Gcc for their experiments. Identifying real occurrences of simultaneous faults is a time-consuming activity, but it can improve the evaluation of SFL techniques.

7 Testing information

The quality of testing information is pivotal for the performance of SFL techniques. Thus, refinements on test suites may impact fault localization performance. A desirable characteristic of test suites is the ability to execute distinguishable parts of the code, which can improve the ability of SFL techniques to more precisely pinpoint faulty code. As large test suites can impact the execution costs of SFL techniques, test suites with reduced size are also desirable. In this section, we present several strategies that propose improvements in testing information for fault localization. Issues regarding evaluation of testing information, coincidental correctness, and use of mutation testing for SFL are also discussed.

7.1 Test suite improvements

Some works have addressed ways to distinguish program entities between test cases to improve fault localization. Baudry et al. [2006] proposed a testing criterion that aims to improve fault localization. They define a structure called *Dynamic Basic Block* (DBB) as a group of statements that are executed by the same test cases. These statements always have the same suspiciousness, and are thereby indistinguishable: the greater the number of DBBs, the lesser the number of indistinguishable statements, and thus the better it is for fault localization. Hao et al. [2010] proposed three strategies to reduce test cases according to their capacity to execute different statements.

Other studies have evaluated the impacts of testing on fault localization and proposed new strategies to improve testing data. Abreu et al. [2007] observed the influence of test suite quality on SFL. They varied the quantity of test cases that exercised faults between passing and failing test cases, and measured the fault localization effectiveness. As expected, having more failing test cases exercising faulty statements leads to better effectiveness. However, a limited number of failing

test cases suffices. In their experiments, a value of six failing runs is optimal, and additional failing test cases have no effect on effectiveness. Xuan and Monperrus [2014] proposed a technique for improving testing information used by SFL techniques. Given a failing test case with more than one assertion, they created one test case for each of these assertions. This aims to avoid situations in which an error occurring in a test case execution prevents the following assertions from being executed. After generating the atomic assertion failing test cases, they apply dynamic slicing to create a list of suspicious statements.

Test suite reduction strategies aim to reduce the amount of test cases keeping the former coverage level. Thus, the execution cost to run the tests reduces, without impacting on the test suite’s quality. These strategies are especially suitable for regression testing. However, reduced test suite size may impact the effectiveness of SFL techniques. Yu et al. [2008] investigated the effects of test suite size reduction on fault localization. They used ten test suite reduction strategies in four ranking metrics. The strategies hold the statement coverage and remove the test cases from different outputs (all test cases, only failing test cases, only passing test cases).

Zhang et al. [2015] used category partition to prioritize test cases for fault localization. Program spectra information is not needed for the prioritization—their technique chooses test cases with inputs farthest from the previous chosen ones, aiming to obtain a high coverage diversity to improve fault localization.

The occurrence of coincidentally correct test cases and their impacts for fault localization have also been studied in the recent years. Coincidental correctness can impair SFL techniques by executing faulty entities as passing test cases, reducing their suspiciousness scores. Masri and Assi [2014] proposed a technique to identify coincidental correct test cases to improve fault localization. They shown that coincidentally correct test cases (CC test cases) are common. They also show that coincidental correctness affects SFL techniques by classifying faulty entities with lesser suspiciousness scores. The proposed technique uses *k-means clustering* to classify test cases as CC or not. Bandyopadhyay and Ghosh [2012] extended the previous version of the work of Masri and Assi [2014], including interactions with the developer to exclude false positive CC test cases. They recalculate the list of remaining suspicious statements throughout the interactions. Other studies that deal with coincidental correctness for SFL have been proposed [Xue et al., 2014, Yang et al., 2015].

Guo et al. [2015] proposed a technique to evaluate the correctness of test oracles. Since humans act as oracles, evaluation mistakes can impair testing and debugging. Their approach considers that tests with similar execution traces likely produce identical results. Similar test cases that diverge are deemed suspicious.

7.2 Mutation testing

Mutation testing has been used to propose new SFL techniques. Nica et al. [2010] proposed a technique to reduce bug candidates by using constraint-based debugging. First, statements that do not violate the constraints and that explain the failing test cases are deemed bug candidates. Second, the technique generates mutants for each bug candidate. Mutants that make the failing test cases pass are used to suggest possible faulty sites. Moon et al. [2014] proposed a technique that uses mutation to modify faulty and correct statements. The rationale is that, if a mutant inserted in a faulty statement reduces the amount of failing test cases, then the faulty statement is more likely to be faulty. Conversely, a mutant inserted in a correct statement which generates more failing test cases is less likely to be faulty. Hong et al. [2015] proposed a similar approach for multilingual programs.

Mutation testing is also used to seed faults for experiments, and to suggest fixes for program repair [Weimer, 2006, Debroy and Wong, 2014]. Ali et al. [2009] used mutation testing to generate

faults and shown that these faults are similar to real faults.

8 Practical use

Spectrum-based Fault Localization’s goal is to help developers to find and fix faults. For practical use, one needs to understand whether the evaluation metrics that assess SFL techniques reflect what happens in development settings. Moreover, the techniques should be assessed by user studies to understand their role in debugging activity.

In this section, we address issues related to the practical use of SFL techniques. First, we present the metrics used to evaluate SFL techniques. We also present experiments with developers using SFL techniques in practice. Finally, we present the strategies proposed to enrich SFL techniques with contextual information.

8.1 Evaluation metrics

There are measures often used by studies to evaluate the performance of SFL techniques. Fault localization effectiveness is an effort measure which indicates how much code is inspected using an SFL technique. As most of the SFL techniques generate ranking lists, studies often use this approach or a variation of it.

The most commonly used metric is *EXAM score* [Chung et al., 2008, Naish et al., 2011, Wong et al., 2012a]. This metric represents the developer’s effort to find a fault using a list of suspicious program entities. The EXAM score is measured as the relative position in which the faulty entity was ranked. It represents the percentage of entities that must be examined to find the fault. EXAM score was based on the metric *score*, proposed by Renieris and Reiss [2003], which indicates the percentage of code that does not need to be examined until finding a fault. Essentially, EXAM score and *score* provide the same information in inversely proportional way. Several works also used *score* [Guo et al., 2006, Ali et al., 2009, Zhao et al., 2011b].

There are other metrics similar to EXAM score. Zhang et al. [2011] proposed a metric called *p-score* to measure the effectiveness of locating suspicious predicates. *Expense* [Yu et al., 2008] is a variation of EXAM score for programs with multiple faults: that is, the percentage of code verified before locating the first fault. To measure the total effort to locate faults for programs with multiple faults, Jones et al. [2007] propose another variation of EXAM score called *total developer expense (D)*, which is the sum of the EXAM score for all faults in a program. Another metric proposed in this work is *critical expense to a failure-free program (FF)*, which measures the time to obtain a failure-free program. Assuming that developers work in parallel to fix the faults, and for each fault found the program is recompiled, FF is the sum of the maximum developer expense at each iteration.

Other metrics identified were *precision* and *recall*, used to measure the accuracy of fault localization techniques based on artificial intelligence. For the fault localization domain, precision generally means the percentage of entities classified by a technique as faults that are in fact faults. Recall is the percentage of faults correctly classified when considering all faults. [Roychowdhury and Khurshid, 2011] used a metric called Metric-Quality to evaluate a technique’s ability to rank the most important statements with higher values, and the least important statements with lower values.

Ranking lists commonly classify program entities with the same suspiciousness scores. This fact impacts the evaluation of ranking metrics, which can vary widely. To deal with ties in ranking lists, Wong et al. [2010] measured the best and the worst cases for the score metric. The best case considers the fault in the first position of the tied entities, while the worst case considers that the

fault is in the last position. Xu et al. [2011] presented a study that shows that ties in SFL ranking lists are common. They propose four tie-breaking strategies to deal with ranking list ties.

Moon et al. [2014] proposed an evaluation metric for fault localization based on *information theory*, called *Locality Information Loss* (LIL). LIL is used to calculate the difference between the true locality and the predicted locality of a fault. This metric can be applied to any technique that generates ranking lists.

Techniques providing lists of suspicious elements often assume a *perfect bug understanding* [Hsu et al., 2008], which supposes that the developer inspecting a list will immediately identify, understand, and fix the fault as soon as s/he reaches the faulty program element. However, this may not happen in practice and the amount of examined code may increase. As pointed out by Parnin and Orso [2011], the measurement of relative positions is quite imprecise. The *absolute number* of entities to be inspected before finding a bug can be a more accurate measure, regardless of the amount of LOC a program has. Liblit et al. [2005] measured the number of predicates their technique returned. Hsu et al. [2008] evaluated their technique of bug signatures (see Subsection 8.3), measuring the absolute number of bug signatures that contain faults. Other studies, most of them from recent years, have used the absolute number of inspected entities to evaluate their techniques [Steimann and Bertschler, 2009, de Souza and Chaim, 2013, Le et al., 2015a]. de Souza and Chaim [2013] also assessed the SFL’s *efficiency*, measuring the number of times a technique found a bug by inspecting less blocks compared to another technique.

The evaluation metrics presented here are useful for comparing SFL techniques in experiments. However, user studies with developers allow us to verify whether these metrics are a good model of what happens in practice.

8.2 User studies

Despite the importance of understanding how SFL techniques can be used in practice, there are few studies that perform experiments with developers. Parnin and Orso [2011] carried out experiments with a group of developers using Tarantula. The authors provided a list of suspicious statements for the developers using two programs, each of them containing a single fault. The results show that the developers take into account their knowledge of the code to search for the faults, and do not usually follow the classification order indicated by the SFL technique. Some other results were observed, including that the perfect bug detection did not occur in the experiment. The authors also verified that the position in which the faulty statement is classified had no significant impact on the ability of developers to find the bugs. The developers suggested improvements, such as the aggregation of the results by their classes or files, and the provision of input values used in test cases to enrich the debugging.

Perez and Abreu [2013] carried out an experiment with 40 developers to assess their technique for visualization of debugging information (GZoltar). The participants were master’s students with more than five years of experience in Java. Two groups were formed—control and experimental groups—with 20 students each using the same program and the same fault. The experimental group used GZoltar and all participants were able to locate the bug in seven minutes on average. The control group used the Eclipse without GZoltar. Only 35% of its participants located the fault within the set time of 30 minutes. Kaleeswaran et al. [2014] carried out a user study to evaluate their program repair technique. Ten developers were asked to work on two independent program repair tasks. For the first task (control phase), each developer received a ranking list with the five most suspicious statements (including the faulty statement) and a test suite of ten passing tests. For the second task (experimental phase) the list with the repair hints was included. In the control phase, eight developers found the fault, and six of them fixed it. In the experimental phase, all the

ten developers found and fixed the bug. They took less time to repair the fault in the experimental phase.

Perscheid et al. [2014] conducted a user study with eight developers to evaluate their program state navigation debugging tool. All developers were undergraduate students with six years of experience. Four faults were debugged by each student, two using the default debugging tool, and two using the new tool. A time limit of fifteen minutes was assigned to each fault. In most cases, the developers found the faults using their approach. They also spent less time to find the same faults by using the new tool.

8.3 Contextual information

SFL techniques have in general tried to precisely pinpoint the faulty site. Ranking lists often contain suspicious elements sorted only by their suspiciousness scores. As a result, elements from different code excerpts can be assigned with higher scores, which may lead to first picks that have no direct relationship among them—for example, statements that do not belong to a same method or class. In practice, when a developer searches for a bug, s/he tries to understand the conditions in which the bug occurs. Techniques have been proposed to support the understanding of the context in which bugs occur. Contextual information in fault localization is associated with strategies that help developers understand bug causes [Jiang and Su, 2007].

Techniques that aim to improve fault localization with contextual information include Jiang and Su [2007], who proposed one of the first techniques for contextualization in fault localization. Their technique selected predicates likely to reveal faults using two machine learning techniques: *Support Vector Machines* (SVM) and *Random Forest* (RF). These predicates are correlated using a *k-means clustering* algorithm. Predicates with similar behaviors over the executions tend to be related. The faulty control-flow paths are constructed based on paths exercised by failing executions that traverse these predicates. The control-flow paths are composed of correlated predicates that provide a context for understanding faults.

Hsu et al. [2008] presented a technique that provides a list of subsequences of elements (branches). These subsequences are called bug signatures; each of them may contain one or more branches in their execution order. The technique first classifies the most suspicious branches. From failing traces, the amount of branches is reduced using a threshold value. They use a *longest common sub-sequence* algorithm to identify subsequences that are present in all failing executions. They are then ordered by their suspicious values. Cheng et al. [2009] extended the idea proposed by [Hsu et al., 2008] using graph mining to present a list of suspicious subgraphs. Graphs of faulty and correct executions are generated to obtain significant subgraphs that differ in the executions. The subgraphs can be extracted at two code levels: blocks or methods.

Hao et al. [2009] proposed an interactive fault localization technique that follows the manual debugging routine. The technique uses the developer’s estimation in the fault localization process. The technique recommends checkpoints based on the suspiciousness of statements. The developer’s feedback is used to update the suspiciousness of statements and choose the next checkpoint.

The technique proposed by R b ler et al. [2012] provides a list of correlated elements likely to be faulty. It combines SFL with automated test generation, using one failing test case and generating several test cases. Only branches and state predicates, called facts, executed in failing test runs are suspected as relevant to faults. Conditional probability is used to estimate the relevance of facts to explain a bug.

Information from source code and code structures is also used to provide contextual information. DiGiuseppe and Jones [2012c] utilize semantic information for fault localization: comments, class and method names, and keywords from the source code. The program is instrumented and the

source code is parsed to extract the information. Terms from the source code are normalized and correlated with the program entities they belong to. A list of top terms is presented as an outcome. de Souza and Chaim [2013] use integration coverage (i.e., pairs of method calls) for SFL. They provide two entity-levels to search for faults. The first level is a list of suspicious methods named roadmap. For each method, it is possible to inspect the most suspicious blocks that belong to it. A threshold is used to limit the number of blocks to be checked for each method, avoiding the inspection of blocks with lesser suspiciousness scores.

Le et al. [2015a] proposed a technique that evaluates the output of SFL techniques (ranking lists) to indicate when this output should be used by developers. They used SVM to build an oracle that indicates whether SFL lists are reliable for inspection. They identified several features of programs to build the oracle, such as number of failing test cases and number of program elements with the same suspiciousness score.

Yi et al. [2015] proposed a technique that combines semantic and dynamic analysis to suggest fault explanations for regression testing. Semantic analysis is applied to identify statements that cause an assertion to fail. Dynamic analysis is then used to identify code changes that retain the failing assertions. These code changes are reported as explanations. The technique presented by Elsaka and Memon [2015] extracts subsequences of statements from a set of failing executions. These subsequences derive from common subsequence graphs and include variable values from the execution.

9 Discussion

In this section, we discuss the main features, results, and challenges of fault localization techniques presented in this survey. We follow the structure proposed in Figure 1 to organize the discussion.

9.1 Techniques

The choice of program spectra influences the performance of SFL techniques, impacting execution costs, data available to be analyzed, and outputs for inspection. Techniques that use coarser spectra data (e.g., method coverage) may have reduced execution costs, requiring less code instrumentation. However, developers will manually inspect these methods, which may increase the amount of code to be verified.

Different spectra have been combined to improve fault localization [Masri, 2010, Santelices et al., 2009, Yu et al., 2011]. We note that data-flow information can improve fault localization techniques. However, the amount of collected data and the execution costs to process it are high compared to those based on control-flow. Strategies to reduce execution costs can make these approaches more feasible for practical use.

Another issue is the assumption of normal distribution of components. Zhang et al. [2011] show that 40% of predicates do not present normal distribution. The author proposed a technique based on non-parametric hypothesis testing for fault localization. This distribution analysis should be investigated for other coverage types, since most of the techniques do not take this issue into account.

Several ranking metrics have been proposed for fault localization, which were created or adapted from other areas. Each has its own specificity. Ochiai differs from Tarantula by taking into account the absence of a statement in failing runs. Jaccard differs from Tarantula by considering statements executed in passing test runs. Experiments have shown that Ochiai have presented higher effectiveness compared to other ranking metrics [Abreu et al., 2007, Le et al., 2013, Xie et al., 2013]. However, the effectiveness of the best ranking metrics is slightly better (e.g., around 1% less code

to examine) in most cases, indicating that they provide similar results [Naish et al., 2011, Le et al., 2013].

Other ways to enhance fault localization explore program behavior; program analysis and artificial intelligence techniques have been used for this purpose. Nevertheless, further studies may be proposed to provide better results.

9.2 Faults

In controlled environments, it may be reasonable to carry out experiments with single-fault programs. However, SFL techniques must deal with an unknown number of faults to be adopted by practitioners. Most of the proposed SFL techniques have been assessed using single-fault programs. Several works have evaluated their performance for multiple faults as a complementary study. These studies carried out complete experiments for programs with single faults, and small experiments with multiple-fault programs [Zhang et al., 2009b, Wong et al., 2012c]. In some cases, experiments for multiple faults use a reduced number of programs when compared to single-fault experiments. SFL techniques have been proposed to deal with multiple faults [Jones et al., 2007, Dean et al., 2009, Xue and Namin, 2013a]. However, their performances have not been compared by these studies.

The absence of benchmarks containing multiple faults makes it difficult to conduct experiments. The studies that perform experiments with multiple faults generate their multiple-fault versions by randomly merging single-fault versions [Jones et al., 2007, Debroy and Wong, 2009, Naish et al., 2009]. This impairs the comparison between techniques, due to the different procedures used to create the benchmarks.

The presence of multiple faults has been shown to hamper fault localization effectiveness [Naish et al., 2009]. On the other hand, DiGiuseppe and Jones [2015] argued that multiple faults had a negligible impact on effectiveness. Their results show that SFL techniques had an average 2% decrease in effectiveness. For large programs, though, 2% of statements may represent a sizable amount of code for inspection.

Another important issue raised by recent studies is the interference between simultaneous faults [Debroy and Wong, 2009, DiGiuseppe and Jones, 2015]. More studies are necessary to investigate the effects of fault interference on fault localization. Existent studies have already shown frequent interference among faults. One issue is that experiments with faults randomly spread across programs do not assure that the faults indeed interfere with each other.

The behavior of SFL techniques in the presence of different fault types is a rarely approached issue. Authors have reported their techniques' difficulties in dealing with particular fault types. There are techniques that explicitly do not identify faults resulting from code omission [Masri, 2010, Yu et al., 2011]. Other authors analyzed the behavior of their techniques for specific faults [Liblit et al., 2005, Mariani et al., 2011, DiGiuseppe and Jones, 2012c]. By assessing the characteristics of faults, one can better understand the strengths and weaknesses of SFL techniques for different fault types.

9.3 Benchmarks

The use of small programs facilitates experiments. It also allows different studies to use the same subject programs to compare techniques. However, the prevalence of small programs or the use of the same benchmark programs in experiments impairs the assessment of SFL techniques. Thus, the use of several programs from different domains is needed for a comprehensive evaluation of SFL techniques.

As discussed in the previous subsection, benchmarks do not contain multiple simultaneous faults. Thus, the creation of benchmark programs with multiple faults can contribute to experimentation in more realistic scenarios. However, creating benchmarks for controlled experimentation is expensive and difficult [Do et al., 2005]. One way to create benchmarks is to identify faults in software repositories of open source programs. This approach provides real faults for experimentation. Another possible solution to reduce the effort in creating benchmarks is to seed faults by using mutation testing, which has been shown to be a good simulation of real faults [Ali et al., 2009].

9.4 Testing information

Software testing is the main source of information for debugging. There are several ways to measure test quality. Testing requirements are used to guarantee that the code is widely tested, and most of the program elements are executed. *Fault detection rate* is a quality measure used to assess the ability of a test suite to reveal failures. Fault localization leads to another desired criterion for test suites: *fault localization capability*. This characteristic means that test suites should be able to distinguish program elements from their test cases.

A natural process for obtaining test suites with higher coverage is to increase their size. However, large test suites lead to greater computational costs to execute them. Test suite reduction strategies are then used to minimize the amount of test cases without losing the ability to failure detection. Test suite reduction techniques are also expected to hold the distinctiveness of program elements throughout their test cases, i.e., the fault localization capability. Thus, test suite reduction techniques have to cope with a trade-off between reducing test size and keeping test cases distinguishable.

9.5 Practical use

Evaluation metrics used to measure SFL techniques are based on assumptions that in practice may not occur. Measuring fault localization performance by the relative position of a faulty entity in a ranking list can mislead the effort to find bugs. For example, if a technique returns the faulty entity within 1% of the most suspicious entities of a program with 100 KLOC, it may be necessary to inspect 1 KLOC to find the fault. This may be infeasible in practice. Developers tend to leave the list if they do not find the fault among the first picks [Parnin and Orso, 2011]. According to Parnin and Orso [2011], the techniques should focus on the absolute position, with the faulty statement among the first positions. Moreover, perfect bug detection (see Subsection 8.1) does not hold in practice, and thus the effort to find faults tends to increase.

The ranking lists provided by SFL techniques are generally based on suspiciousness scores, and may be composed of entities with no direct relation to each other (e.g., statements that belong to distinct classes). This fact may impair the identification of faults. Thus, techniques that provides more information can help debuggers understand the conditions in which faults occur. Strategies have been proposed to tackle this issue, such as grouping related entities, exploring different code levels, adding source code information, or presenting subsequences of execution traces. Future approaches must explore new ways to reduce the amount of non-relevant information and to enrich the information provided to developers.

When developing fault localization techniques, we suppose several assumptions about the developers' behavior while performing debugging tasks. However, without user studies, one cannot know whether these assumptions hold in practice. Thus, these studies are essential for identifying how the techniques are used, to assess the developers' fondness of SFL techniques, and to provide guidance on their use in industrial settings. Unfortunately, debugging user studies are still scarce,

though they are pivotal for leading to SFL adoption by practitioners.

9.6 Concept map in Spectrum-based Fault Localization

Concept maps [Novak and Cañas, 2008] are used to organize and represent knowledge. Based on the analysis process carried out throughout this survey, we created a concept map representing the main concepts and their relationships regarding the spectrum-based fault localization area.

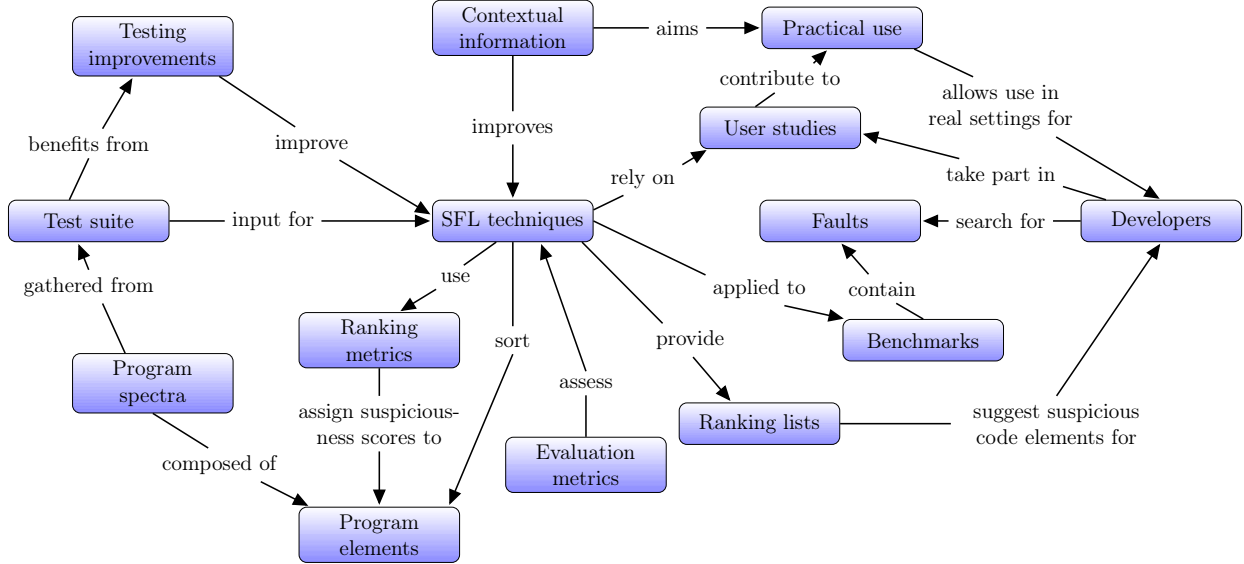


Figure 2: Concept map on Spectrum-based Fault Localization

10 Related work

Other studies were proposed to provide an overview of the fault localization area. Alipour [2011] conducted a survey on fault localization. The author considered that the major fault localization approaches are program slicing, spectrum-based, statistical inference, delta debugging, dynamic, and model checking. In his survey, only six studies of SFL techniques were addressed—most of the studies are model checking-based techniques. The author concludes that such techniques are far from practical use due to issues related to execution time and scalability for large programs. Beyond this concerns, model checking techniques usually requires formal specifications of programs, which is difficult to obtain for most programs. Agarwal and Agrawal [2014] presented a literature review on fault localization, including studies from 2007 to 2013. They selected 30 papers from main Software Engineering journals and conferences. Most of the papers are focused on test suite improvements for fault localization and SFL techniques. The results are presented in a table describing studies’ characteristics and a description of the most frequent techniques and strategies in the area.

[Wong et al., 2016] presented a fault localization survey addressing techniques from 1977 to November 2014. They classified the techniques in eight categories: program slicing, spectrum-based, statistics, program state, machine learning, data mining, model-based debugging, and additional techniques. This work also addresses fault localization tools developed by the studies presented.

Our work differs by addressing spectrum-based techniques from 2005 to February 2016, although we also discussed seminal works from 1950s to 2004 through a historical overview. Besides the

database search, we applied a snowballing process to extend the searching for fault localization studies. We included studies that focus on testing improvements for fault localization, and also mutation-based techniques. Moreover, we addressed issues related to the practical adoption of SFL, such as user studies and techniques that provide additional contextual information, aiming to improve developers' program comprehension. We understand that future research must focus on strategies to allow the use of SFL techniques in real settings. Another contribution of this survey is to propose a concept map regarding the main concepts of SFL techniques.

11 Conclusion

A great number of fault localization techniques have been proposed in the last decades. These techniques aim to pinpoint program entities that are likely to be faulty. Thus, developers can inspect these entities to find faults, reducing the time spent debugging.

This survey focuses on spectrum-based fault localization (SFL) techniques, which have presented promising results. We address the main issues regarding SFL to provide a comprehensive overview about the research area: SFL techniques, faults, benchmarks, testing information, and practical use.

Several advances have been achieved, whilst a number of challenges and limitations should be tackled to improve SFL techniques. New ways for exploring program spectrum information and new strategies for generating reduced sets of suspicious entities can contribute to improving the output results. Combining different spectra (e.g., data-flow and control-flow spectra) seems to fine tune the fault localization ability of SFL techniques; however, sophisticated spectra are costly to collect. Strategies for collecting fine-grained coverage levels from suspicious coarser levels can help balance execution costs and output precision. New techniques that cope with multiple-fault programs are needed to support fault localization of real programs, in which the number of faults is unknown. Large-size benchmark programs, using different fault types and multiple faults, will provide realistic scenarios for assessing SFL techniques. More user studies will enable better understanding of how fault localization techniques are used in practice. This survey also presents a concept map of SFL, representing the relationships between the main topics and challenges for future research. By presenting the state-of-the-art of SFL techniques, we hope this survey encourages the development of debugging techniques that end up adopted by practitioners.

Acknowledgments

This work is supported by FAPESP (Fundação de Amparo à Pesquisa do Estado de São Paulo – São Paulo Research Foundation), under grants 2013/24992-2 and 2014/23030-5.

References

- Rui Abreu, Peter Zoetewij, and Arjan J. C. van Gemund. On the accuracy of spectrum-based fault localization. In *Proceedings of the Testing: Academic and Industrial Conference Practice and Research Techniques - MUTATION*, TAICPART-MUTATION'07, pages 89–98, 2007.
- Rui Abreu, Peter Zoetewij, and Arjan J. C. van Gemund. Simultaneous debugging of software faults. *Journal of Systems and Software*, 84(4):573–586, 2011.
- Anne Adam and Jean-Pierre Laurent. Laura, a system to debug student programs. *Artificial Intelligence*, 15(1–2):75–122, 1980.

- Pragya Agarwal and Arun P. Agrawal. Fault-localization techniques for software systems: A literature review. *ACM SIGSOFT Software Engineering Notes*, 39(5):1–8, 2014.
- Hiralal Agrawal and Eugene H. Spafford. Bibliography on debugging and backtracking. *ACM SIGSOFT Software Engineering Notes*, 14(2):49–56, 1989.
- Hiralal Agrawal, Joseph R. Horgan, S. London, and W. Eric Wong. Fault localization using execution slices and dataflow tests. In *Proceedings of the 6th IEEE International Symposium on Software Reliability Engineering*, ISSRE’95, pages 143–151, 1995.
- Shaimaa Ali, James H. Andrews, Tamilselvi Dhandapani, and Wantao Wang. Evaluating the accuracy of fault localization techniques. In *Proceedings of the 24th IEEE/ACM International Conference on Automated Software Engineering*, ASE’09, pages 76–87, 2009.
- Mohammad A. Alipour. Automated fault localization techniques; a survey, 2011. URL <http://web.engr.oregonstate.edu/~alipour/pub/flsurvey2.pdf>.
- Elton Alves, Milos Gligoric, Vilas Jagannath, and Marcelo d’Amorim. Fault-localization using dynamic slicing and change impact analysis. In *Proceedings of the 26th IEEE/ACM International Conference on Automated Software Engineering*, ASE’11, pages 520–523, 2011.
- Keijiro Araki, Zengo Furukawa, and Jingde Cheng. A general framework for debbugging. *IEEE Software*, 8(3):14–20, 1991.
- Aristotle Research Group. Aristotle analysis system, 2007. URL <https://research.cc.gatech.edu/aristotle>.
- George K. Baah, Andy Podgurski, and Mary J. Harrold. Causal inference for statistical fault localization. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA’10, pages 73–84, 2010.
- R. M. Balzer. Exdams: extendable debugging and monitoring system. In *Proceedings of the Spring Joint Computer Conference*, AFIPS’69 (Spring), pages 567–580, 1969.
- Aritra Bandyopadhyay and Sudipta Ghosh. Tester feedback driven fault localization. In *Proceedings of the 5th IEEE International Conference on Software Testing, Verification and Validation*, ICST’12, pages 41–50, 2012.
- Aritra Bandyopadhyay and Sudipto Ghosh. Proximity based weighting of test cases to improve spectrum based fault localization. In *Proceedings of the 26th IEEE/ACM International Conference on Automated Software Engineering*, ASE’11, pages 420–423, 2011.
- Benoit Baudry, Franck Fleurey, and Yves Le Traon. Improving test suites for efficient fault localization. In *Proceedings of the 28th International Conference on Software Engineering*, ICSE’06, pages 82–91, 2006.
- Martin Burger and Andreas Zeller. Minimizing reproduction of software failures. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA’11, pages 221–231, 2011.
- Hong Cai and Xing Xu. A new spectrum-based fault localization method by using clustering algorithm. *International Journal of Advancements in Computing Technology*, 4(22):848–856, 2012.

- Mike Y. Chen, Emre Kiciman, Eugene Fratkin, Armando Fox, and Eric Brewer. Pinpoint: Problem determination in large, dynamic internet services. In *Proceedings of the 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, DSN'02, pages 595–604, 2002.
- Hong Cheng, David Lo, Yang Zhou, Xiaoyin Wang, and Xifeng Yan. Identifying bug signatures using discriminative graph mining. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'09, pages 141–152, 2009.
- Trishul M. Chilimbi, Ben Liblit, Krishna Mehra, Aditya V. Nori, and Kapil Vaswani. Holmes: Effective statistical debugging via efficient path profiling. In *Proceedings of the 31st International Conference on Software Engineering*, ICSE'09, pages 34–44, 2009.
- Yu-Min Chung, Chin-Yu Huang, and Yu-Chi Huang. A study of modified testing-based fault localization method. In *Proceedings of the 14th IEEE Pacific Rim International Symposium on Dependable Computing*, PRDC'08, pages 168–175, 2008.
- James S. Collofello and Larry Cousins. Toward automatic software fault localization through decision-to-decision path analysis. In *Proceedings of the AFIP 1987 National Computer Conference*, pages 539–544, 1987.
- Valentin Dallmeier, Christian Lindig, and Andreas Zeller. Lightweight defect localization for java. In *Proceedings of the 19th European Conference on Object-Oriented Programming*, ECOOP'05, pages 528–550, 2005.
- Gong Dandan, Su Xiaohong, Wang Tiantian, Ma Peijun, and Yu Wang. State dependency probabilistic model for fault localization. *Information and Software Technology*, 57(0):430–445, 2014.
- Higor A. de Souza and Marcos L. Chaim. Adding context to fault localization with integration coverage. In *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering*, ASE'13, pages 628–633, 2013.
- Brian C. Dean, William B. Pressly, Brian A. Malloy, and Adam A. Whitley. A linear programming approach for automated localization of multiple faults. In *Proceedings of the 24th IEEE/ACM International Conference on Automated Software Engineering*, ASE'09, pages 640–644, 2009.
- Vidroha Debroy and W. Eric Wong. Insights on fault interference for programs with multiple bugs. In *Proceedings of the 20th IEEE International Symposium on Software Reliability Engineering*, ISSRE'09, pages 165–174, 2009.
- Vidroha Debroy and W. Eric Wong. On the equivalence of certain fault localization techniques. In *Proceedings of the 26th ACM Symposium on Applied Computing*, SAC'11, pages 1457–1463, 2011a.
- Vidroha Debroy and W. Eric Wong. On the consensus-based application of fault localization techniques. In *Proceedings of the 2011 IEEE 35th Annual Computer Software and Applications Conference Workshops*, COMPSACW'11, pages 506–511, 2011b.
- Vidroha Debroy and W. Eric Wong. Combining mutation and fault localization for automated program debugging. *Journal of Systems and Software*, 90(0):45–60, 2014.
- Vidroha Debroy, W. Eric Wong, Xiaofeng Xu, and Byoungju Choi. A grouping-based strategy to improve the effectiveness of fault localization techniques. In *Proceedings of the 10th International Conference on Quality Software*, QSIC'10, pages 13–22, 2010.

- Nicholas DiGiuseppe and James A. Jones. Software behavior and failure clustering: An empirical study of fault causality. In *Proceedings of the 5th IEEE International Conference on Software Testing, Verification and Validation*, ICST'12, pages 191–200, 2012a.
- Nicholas DiGiuseppe and James A. Jones. Concept-based failure clustering. In *Proceedings of the 20th ACM SIGSOFT International Symposium on the Foundations of Software Engineering*, FSE'12, pages 29:1–29:4, 2012b.
- Nicholas DiGiuseppe and James A. Jones. Semantic fault diagnosis: automatic natural-language fault descriptions. In *Proceedings of the 20th ACM SIGSOFT International Symposium on the Foundations of Software Engineering*, FSE'12, pages 23:1–23:4, 2012c.
- Nicholas DiGiuseppe and James A. Jones. Fault density, fault types, and spectra-based fault localization. *Empirical Software Engineering*, 20(4):928–967, 2015.
- Hyunsook Do, Sebastian Elbaum, and Gregg Rothermel. Supporting controlled experimentation with testing techniques: An infrastructure and its potential impact. *Empirical Software Engineering*, 10(4):405–435, 2005.
- João Durães and Henrique Madeira. Emulation of software faults: A field data study and a practical approach. *IEEE Transactions on Software Engineering*, 32(11):849–867, 2006.
- Ethar Elsaka and Atif Memon. Disqover: Debugging via code sequence covers. In *Proceedings of the IEEE 26th International Symposium on Software Reliability Engineering Workshops*, ISSREW'15, pages 85–92, 2015.
- Thomas G. Evans and D. Lucille Darley. On-line debugging techniques: A survey. In *Proceedings of the Fall Joint Computer Conference*, AFIPS'66 (Fall), pages 37–50, 1966.
- Richard E. Fairley. Aladdin: Assembly language assertion driven debugging interpreter. *IEEE Transactions on Software Engineering*, 5(4):426–428, 1979.
- Peter Fritzson, Nahid Shahmehri, Mariam Kamkar, and Tibor Gyimothy. Generalized algorithmic debugging and testing. *ACM Letters on Programming Languages and Systems*, 1(4):303–322, 1992.
- John T. Gilmore. Tx-o direct input utility system. Memo 6M-5097, 1957. Lincoln Laboratory, MIT.
- Liang Guo, Abhik Roychoudhury, and Tao Wang. Accurately choosing execution runs for software fault localization. In Alan Mycroft and Andreas Zeller, editors, *Compiler Construction*, volume 3923 of *Lecture Notes in Computer Science*, pages 80–95. Springer, Berlin, 2006.
- Xinrui Guo, Min Zhou, Xiaoyu Song, Ming Gu, and Jianguang Sun. First, debug the test oracle. *IEEE Transactions on Software Engineering*, 41(10):986–1000, 2015.
- Dan Hao, Lu Zhang, Tao Xie, Hong Mei, and Jia-Su Sun. Interactive fault localization using test information. *Journal of Computer Science and Technology*, 24(5):962–974, 2009.
- Dan Hao, Tao Xie, Lu Zhang, Xiaoyin Wang, Jiasu Sun, and Hong Mei. Test input reduction for result inspection to facilitate fault localization. *Automated Software Engineering*, 17(1):5–31, 2010.

- Wolfgang Högerle, Friedrich Steimann, and Marcus Frenkel. More debugging in parallel. In *Proceedings of the 25th IEEE International Symposium on Software Reliability Engineering*, ISSRE'14, pages 133–143, 2014.
- Shin Hong, Byeongcheol Lee, Taehoon Kwak, Yiru Jeon, Bongsuk Ko, Yunho Kim, and Moonzoo Kim. Mutation-based fault localization for real-world multilingual programs (t). In *Proceedings of the 30th IEEE/ACM International Conference on Automated Software Engineering*, ASE'15, pages 464–475, 2015.
- David Hovemeyer and William Pugh. Finding bugs is easy. *SIGPLAN Notices*, 39(12):92–106, 2004.
- Hwa-You Hsu, James A. Jones, and Alessandro Orso. Rapid: Identifying bug signatures to support debugging activities. In *Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering*, ASE'08, pages 439–442, 2008.
- Monica Hutchins, Herb Foster, Tarak Goradia, and Thomas Ostrand. Experiments of the effectiveness of dataflow- and controlflow-based test adequacy criteria. In *Proceedings of the 16th International Conference on Software Engineering*, ICSE'94, pages 191–200, 1994.
- IEEE. Ieee standard glossary of software engineering terminology (ieee std 610.12-1990), 1990.
- Samireh Jalali and Claes Wohlin. Systematic literature studies: Database searches vs. backward snowballing. In *Proceedings of the 6th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ESEM'12, pages 29–38, 2012.
- Dennis Jeffrey, Neelam Gupta, and Rajiv Gupta. Fault localization using value replacement. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'08, pages 167–178, 2008.
- Lingxiao Jiang and Zhendong Su. Context-aware statistical debugging: From bug predictors to faulty control flow paths. In *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering*, ASE'07, pages 184–193, 2007.
- Wei Jin and Alessandro Orso. F3: Fault localization for field failures. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'13, pages 213–223, 2013.
- W. Lewis Johnson and Elliot Soloway. Proust: Knowledge-based program understanding. *IEEE Transactions on Software Engineering*, 11(3):267–275, 1985.
- James A. Jones, Mary J. Harrold, and John Stasko. Visualization of test information to assist fault localization. In *Proceedings of the 24th International Conference on Software Engineering*, ICSE'02, pages 467–477, 2002.
- James A. Jones, James F. Bowring, and Mary J. Harrold. Debugging in parallel. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'07, pages 16–26, 2007.
- Xiaolin Ju, Shujuan Jiang, Xiang Chen, Xingya Wang, Yanmei Zhang, and Heling Cao. Hsfal: Effective fault localization using hybrid spectrum of full slices and execution slices. *Journal of Systems and Software*, 90(0):3–17, 2013.

- Shalini Kaleeswaran, Varun Tulsian, Aditya Kanade, and Alessandro Orso. Minthint: Automated synthesis of repair hints. In *Proceedings of the 36th International Conference on Software Engineering*, ICSE'14, pages 266–276, 2014.
- Peggy A. Kidwell. Stalking the elusive computer bug. *IEEE Annals of the History of Computing*, 20(4):5–9, 1998.
- Jeongho Kim and Eunseok Lee. Empirical evaluation of existing algorithms of spectrum based fault localization. In *Proceedings of the 28th International Conference on Information Networking*, ICOIN'14, pages 346–351, 2014.
- Bogdan Korel. Pelas-program error-locating assistant system. *IEEE Transactions on Software Engineering*, 14(9):1253–1260, 1988.
- Bogdan Korel and Janus Laski. Dynamic program slicing. *Information Processing Letters*, 29(3):155–163, 1988.
- Alan Kotok. Dec debugging tape. Memo MIT-1, August 1961. MIT.
- Gulsher Laghari, Alessandro Murgia, and Serge Demeyer. Localising faults in test execution traces. In *Proceedings of the 14th International Workshop on Principles of Software Evolution*, IWPSE 2015, pages 1–8, 2015.
- Tien-Duy B. Le, Ferdian Thung, and David Lo. Theory and practice, do they match? a case with spectrum-based fault localization. In *Proceedings of the 29th IEEE International Conference on Software Maintenance*, ICSM'13, pages 380–383, 2013.
- Tien-Duy B. Le, David Lo, and Ferdian Thung. Should i follow this fault localization tool's output? *Empirical Software Engineering*, 20(5):1237–1274, 2015a.
- Tien-Duy B. Le, Richard J. Oentaryo, and David Lo. Information retrieval and spectrum based bug localization: Better together. In *Proceedings of the 10th Joint Meeting of the European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering*, ESEC/FSE 2015, pages 579–590, 2015b.
- Heng Li, Yuzhen Liu, Zhenyu Zhang, and Jian Liu. Program structure aware fault localization. In *Proceedings of the International Workshop on Innovative Software Development Methodologies and Practices*, InnoSWDev'14, pages 40–48, 2014.
- Ben Liblit, Mayur Naik, Alice X. Zheng, Alex Aiken, and Michael I. Jordan. Scalable statistical bug isolation. In *Proceedings of the 26th Annual ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI'05, pages 15–26, 2005.
- Chao Liu, Xifeng Yan, Hwanjo Yu, Jiawei Han, and Philip S. Yu. Mining behavior graphs for "backtrace" of noncrashing bugs. In *Proceedings of the 2005 SIAM International Conference on Data Mining*, SDM'05, pages 286–297, 2005.
- Chao Liu, Long Fei, X. Yan, Jiawei Han, and Samuel P. Midkiff. Statistical debugging: A hypothesis testing-based approach. *IEEE Transactions on Software Engineering*, 32(10):831–848, 2006.
- Chao Liu, Xiangyu Zhang, and Jiawei Han. A systematic study of failure proximity. *IEEE Transactions on Software Engineering*, 34(6):826–843, 2008.

- Xuemei Liu, Yongpo Liu, Ji Wu, and Xiaoxia Jia. Finding suspicious patterns of object-oriented programs based on variance analysis. In *Proceedings of the 2010 7th International Conference on Fuzzy Systems and Knowledge Discovery*, FSKD'10, pages 2815–2820, 2010.
- Lucia Lucia, David Lo, Lingxiao Jiang, Ferdian Thung, and Aditya Budi. Extended comprehensive study of association measures for fault localization. *Journal of Software: Evolution and Process*, 26(2):172–219, 2014.
- Chunyan Ma, Yifei Zhang, Tao Zhang, Yuwei Lu, and Qingyi Wang. Uniformly evaluating and comparing ranking metrics for spectral fault localization. In *Proceedings of the 14th International Conference on Quality Software*, QSIC'14, pages 315–320, 2014.
- Leonardo Mariani, Fabrizio Pastore, and Mauro Pezze. Dynamic analysis for diagnosing integration faults. *IEEE Transactions on Software Engineering*, 37(4):486–508, 2011.
- Wes Masri. Fault localization based on information flow coverage. *Software Testing, Verification and Reliability*, 20(2):121–147, 2010.
- Wes Masri and Rawad A. Assi. Prevalence of coincidental correctness and mitigation of its impact on fault localization. *ACM Transactions on Software Engineering and Methodology*, 23(1):8:1–8:28, 2014.
- Seokhyeon Moon, Yunho Kim, Moonzoo Kim, and Shin Yoo. Ask the mutants: Mutating faulty programs for fault localization. In *Proceedings of the 7th IEEE International Conference on Software Testing, Verification and Validation*, ICST'14, pages 153–162, 2014.
- Syed S. Murtaza, Mechelle Gittens, and Nazim H. Madhavji. Discovering the fault origin from field traces. In *Proceedings of the 19th IEEE International Symposium on Software Reliability Engineering*, ISSRE'08, pages 295–296, 2008.
- Glenford J. Myers. *The Art of Software Testing*. John Wiley & Sons, Inc., New York, NY, 1979.
- George Nagy and M. Carlson Pennebaker. A step toward automatic analysis of student programming errors in a batch environment. *International Journal of Man-Machine Studies*, 6(5):563–578, 1974.
- Piramanayagam A. Nainar, Ting Chen, Jake Rosin, and Ben Liblit. Statistical debugging using compound boolean predicates. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'07, pages 5–15, 2007.
- Lee Naish and Hua Jie Lee. Duals in spectral fault localization. In *Proceedings of the 22nd Australian Software Engineering Conference*, ASWEC'13, pages 51–59, 2013.
- Lee Naish, Hua Jie Lee, and Kotagiri Ramamohanarao. Spectral debugging with weights and incremental ranking. In *Proceedings of the 16th Asia-Pacific Software Engineering Conference*, APSEC'09, pages 168–175, 2009.
- Lee Naish, Hua Jie Lee, and Kotagiri Ramamohanarao. Statements versus predicates in spectral bug localization. In *Proceedings of the 17th Asia-Pacific Software Engineering Conference*, APSEC'10, pages 375–384, 2010.
- Lee Naish, Hua Jie Lee, and Kotagiri Ramamohanarao. A model for spectra-based software diagnosis. *ACM Transactions on Software Engineering and Methodology*, 20(3):1–32, 2011.

- Syeda Nessa, Muhammad Abedin, W. Eric Wong, Latifur Khan, and Yu Qi. Software fault localization using n-gram analysis. In *Proceedings of the 3rd International Conference on Wireless Algorithms, Systems, and Applications*, WASA'08, pages 548–559, 2008.
- Mihai Nica, Simona Nica, and Franz Wotawa. Does testing help to reduce the number of potentially faulty statements in debugging? In Leonardo Bottaci and Gordon Fraser, editors, *Testing - Practice and Research Techniques*, volume 6303 of *Lecture Notes in Computer Science*, pages 88–103. Springer, Berlin, 2010.
- Joseph D Novak and Alberto J Cañas. The theory underlying concept maps and how to construct and use them. Technical report, Institute for Human and Machine Cognition, Pensacola, United States, 2008.
- Mike Papadakis and Yves Le Traon. Metallaxis-fl: Mutation-based fault localization. *Software Testing, Verification and Reliability*, 2013.
- Chris Parnin and Alessandro Orso. Are automated debugging techniques actually helping programmers? In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'11, pages 199–209, 2011.
- Alexandre Perez and Rui Abreu. Cues for scent intensification in debugging. In *Proceedings of the IEEE 24th International Symposium on Software Reliability Engineering Workshops*, ISSREW'13, pages 120–125, 2013.
- Alexandre Perez, André Riboira, and Rui Abreu. A topology-based model for estimating the diagnostic efficiency of statistics-based approaches. In *Proceedings of the IEEE 23rd International Symposium on Software Reliability Engineering Workshops*, ISSREW'12, pages 171–176, 2012.
- Michael Perscheid, Tim Felgentreff, and Robert Hirschfeld. Follow the path: Debugging state anomalies along execution histories. In *Proceedings of the 2014 Software Evolution Week - IEEE Conference on Software Maintenance, Reengineering and Reverse Engineering*, CSMR-WCRE'14, pages 124–133, 2014.
- Manos Renieris and Steven P. Reiss. Fault localization with nearest neighbor queries. In *Proceedings of the 18th IEEE International Conference on Automated Software Engineering*, ASE'03, pages 30–39, 2003.
- Thomas Reps, Thomas Ball, Manuvir Das, and James Larus. The use of program profiling for software maintenance with applications to the year 2000 problem. In *Proceedings of the 6th European Software Engineering Conference Held Jointly with the 5th ACM SIGSOFT Symposium on the Foundations of Software Engineering*, ESEC/FSE'97, pages 432–449, 1997.
- Jeremias Röbler, Gordon Fraser, Andreas Zeller, and Alessandro Orso. Isolating failure causes through test case generation. In *Proceedings of the ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA'12, pages 309–319, 2012.
- Shounak Roychowdhury. Ensemble of feature selectors for software fault localization. In *Proceedings of the 2012 IEEE International Conference on Systems, Man, and Cybernetics*, SMC'12, pages 1351–1356, 2012.
- Shounak Roychowdhury and Sarfraz Khurshid. Software fault localization using feature selection. In *Proceedings of the International Workshop on Machine Learning Technologies in Software Engineering*, MALETS'11, pages 11–18, 2011.

- Nick Rutar, Christian B. Almazan, and Jeffrey S. Foster. A comparison of bug finding tools for java. In *Proceedings of the 15th IEEE International Symposium on Software Reliability Engineering, ISSRE'04*, pages 245–256, 2004.
- Raul Santelices, James A. Jones, Yanbing Yu, and Mary J. Harrold. Lightweight fault-localization using multiple coverage types. In *Proceedings of the 31st International Conference on Software Engineering, ICSE'09*, pages 56–66, 2009.
- Ehud Y. Shapiro. *Algorithmic program debugging*. MIT Press, Cambridge, MA, 1983.
- Friedrich Steimann and Mario Bertschler. A simple coverage-based locator for multiple faults. In *Proceedings of the 2nd IEEE International Conference on Software Testing, Verification and Validation, ICST'09*, pages 366–375, 2009.
- Thomas G. Stockham and Jack B. Dennis. Flit - flexowriter interrogation tape: A symbolic utility program for tx-o. Memo 5001-23, July 1960. Department of Electric Engineering, MIT.
- Gregory Tasse. The economic impacts of inadequate infrastructure for software testing. *National Institute of Standards and Technology, RTI Project, 7007(011)*, 2002.
- Ferdian Thung, David Lo, and Lingxiao Jiang. Automatic defect categorization. In *Proceedings of the 19th Working Conference on Reverse Engineering, WCRE'12*, pages 205–214, 2012.
- Wang Tiantian, Su Xiaohong, Ma Peijun, and Wang Kechao. Comprehension oriented software fault location. In *Proceedings of the 2011 International Conference on Computer Science and Network Technology, ICCSNT'11*, pages 340–343, 2011.
- Filippos I. Vokolos and Phyllis G. Frankl. Empirical evaluation of the textual differencing regression testing technique. In *Proceedings of the 16th IEEE International Conference on Software Maintenance, ICSM'98*, pages 44–53, 1998.
- Xiaomin Wan, Xiaoguang Mao, and Ziyang Dai. Fault localization via behavioral models. In *Proceedings of the 3rd International Conference on Software Engineering and Service Science, ICSESS'12*, pages 472–475, 2012.
- Tao Wang and Abhik Roychoudhury. Automated path generation for software fault localization. In *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering, ASE'05*, pages 347–351, 2005.
- Xiaoyan Wang and Yongmei Liu. Automated fault localization via hierarchical multiple predicate switching. *Journal of Systems and Software*, 104(0):69–81, 2015.
- Xinping Wang, Qing Gu, Xin Zhang, Xiang Chen, and Daoxu Chen. Fault localization based on multi-level similarity of execution traces. In *Proceedings of the 16th Asia-Pacific Software Engineering Conference, APSEC'09*, pages 399–405, 2009.
- Westley Weimer. Patches as better bug reports. In *Proceedings of the 5th International Conference on Generative Programming and Component Engineering, GPCE'06*, pages 181–190, 2006.
- Mark Weiser. Program slicing. In *Proceedings of the 5th International Conference on Software Engineering, ICSE'81*, pages 439–449, 1981.

- Wanzhi Wen, Bixin Li, Xiaobing Sun, and Jiakai Li. Program slicing spectrum-based software fault localization. In *Proceedings of the 23rd International Conference on Software Engineering and Knowledge Engineering, SEKE'11*, pages 213–218, 2011.
- W. Eric Wong and Yu Qi. Effective program debugging based on execution slices and inter-block data dependency. *Journal of Systems and Software*, 79(7):891–903, 2006.
- W. Eric Wong, Vidroha Debroy, and Byoungju Choi. A family of code coverage-based heuristics for effective fault localization. *Journal of Systems and Software*, 83(2):188–208, 2010.
- W. Eric Wong, Vidroha Debroy, Richard Golden, Xiaofeng Xu, and Bhavani Thuraisingham. Effective software fault localization using an rbf neural network. *IEEE Transactions on Reliability*, 61(1):149–169, 2012a.
- W. Eric Wong, Vidroha Debroy, Yihao Li, and Ruizhi Gao. Software fault localization using dstar (d*). In *Proceedings of the IEEE 6th International Conference on Software Security and Reliability, SERE'12*, pages 21–30, 2012b.
- W. Eric Wong, Vidroha Debroy, and Dianxiang Xu. Towards better fault localization: A crosstab-based statistical approach. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 42(3):378–396, 2012c.
- W. Eric Wong, Ruizhi Gao, Yihao Li, Rui Abreu, and Franz Wotawa. A survey on software fault localization. *IEEE Transactions on Software Engineering*, 2016. To appear.
- Franz Wotawa, Markus Stumptner, and Wolfgang Mayer. Model-based debugging or how to diagnose programs automatically. In Tim Hendtlass and Moonis Ali, editors, *Developments in Applied Artificial Intelligence*, volume 2358 of *Lecture Notes in Computer Science*, pages 243–257. Springer, Berlin, 2002.
- Xiaoyuan Xie, Tsong Yueh Chen, and Baowen Xu. Isolating suspiciousness from spectrum-based fault localization techniques. In *Proceedings of the 10th International Conference on Quality Software, QSIC'10*, pages 385–392, 2010.
- Xiaoyuan Xie, Tsong Yueh Chen, Fei-Ching Kuo, and Baowen Xu. A theoretical analysis of the risk evaluation formulas for spectrum-based fault localization. *ACM Transactions on Software Engineering and Methodology*, 22(4):31:1–31:40, 2013.
- Jian Xu, Zhenyu Zhang, W. K. Chan, T. H. Tse, and Shanping Li. A general noise-reduction framework for fault localization of java programs. *Information and Software Technology*, 55(5):880–896, 2013.
- Xiaofeng Xu, Vidroha Debroy, W. Eric Wong, and Donghui Guo. Ties within fault localization rankings: Exposing and addressing the problem. *International Journal of Software Engineering and Knowledge Engineering*, 21(6):803–827, 2011.
- Jifeng Xuan and Martin Monperrus. Test case purification for improving fault localization. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on the Foundations of Software Engineering, FSE'14*, pages 52–63, 2014.
- Xiaozhen Xue and Akbar S. Namin. Measuring the odds of statements being faulty. In Hubert B. Keller, Erhard Plodereder, Peter Dencker, and Herbert Klenk, editors, *Reliable Software Technologies - Ada-Europe 2013*, volume 7896 of *Lecture Notes in Computer Science*, pages 109–126. Springer, Berlin, 2013a.

- Xiaozhen Xue and Akbar S. Namin. How significant is the effect of fault interactions on coverage-based fault localizations? In *Proceedings of the 7th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ESEM'13, pages 113–122, 2013b.
- Xiaozhen Xue, Yulei Pang, and Akbar S. Namin. Trimming test suites with coincidentally correct test cases for enhancing fault localizations. In *Proceedings of the IEEE 38th Annual International Computers, Software and Applications Conference*, COMPSAC'14, pages 239–244, 2014.
- Xiaoshuang Yang, Mengleng Liu, Ming Cao, Lei Zhao, and Lina Wang. Regression identification of coincidental correctness via weighted clustering. In *Proceedings of the IEEE 39th Annual International Computers, Software and Applications Conference*, COMPSAC'15, pages 115–120, 2015.
- Qiuping Yi, Zijiang Yang, Jian Liu, Chen Zhao, and Chao Wang. A synergistic analysis method for explaining failed regression tests. In *Proceedings of the 37th International Conference on Software Engineering*, ICSE'15, pages 257–267, 2015.
- Cemal Yilmaz, Amit Paradkar, and Clay Williams. Time will tell: Fault localization using time spectra. In *Proceedings of the 30th International Conference on Software Engineering*, ICSE'08, pages 81–90, 2008.
- Shin Yoo. Evolving human competitive spectra-based fault localisation techniques. In Gordon Fraser and Jefferson Teixeira de Souza, editors, *Search Based Software Engineering - SSBSE 2012*, volume 7515 of *Lecture Notes in Computer Science*, pages 244–258. Springer, Berlin, 2012.
- Kai Yu, Mengxiang Lin, Qing Gao, Hui Zhang, and Xiangyu Zhang. Locating faults using multiple spectra-specific models. In *Proceedings of the 26th ACM Symposium on Applied Computing*, SAC'11, pages 1404–1410, 2011.
- Yanbing Yu, James A. Jones, and Mary J. Harrold. An empirical study of the effects of test-suite reduction on fault localization. In *Proceedings of the 30th International Conference on Software Engineering*, ICSE'08, pages 201–210, 2008.
- Zhongxing Yu, Chenggang Bai, and Kai-Yuan Cai. Does the failing test execute a single or multiple faults?: An approach to classifying failing tests. In *Proceedings of the 37th International Conference on Software Engineering*, ICSE'15, pages 924–935, 2015.
- Andreas Zeller. Isolating cause-effect chains from computer programs. In *Proceedings of the 10th ACM SIGSOFT International Symposium on the Foundations of Software Engineering*, FSE'02, pages 1–10, 2002.
- Sai Zhang and Congle Zhang. Software bug localization with markov logic. In *Proceedings of the 36th International Conference on Software Engineering*, ICSE'14, pages 424–427, 2014.
- Xiangyu Zhang, Haifeng He, Neelam Gupta, and Rajiv Gupta. Experimental evaluation of using dynamic slices for fault location. In *Proceedings of the 6th International Symposium on Automated Analysis-driven Debugging*, AADeBUG'05, pages 33–42, 2005.
- Xiangyu Zhang, Neelam Gupta, and Rajiv Gupta. Locating faults through automated predicate switching. In *Proceedings of the 28th International Conference on Software Engineering*, ICSE'06, pages 272–281, 2006.

- Xiangyu Zhang, Sriraman Tallam, Neelam Gupta, and Rajiv Gupta. Towards locating execution omission errors. In *Proceedings of the 28th Annual ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI'07, pages 415–424, 2007.
- Xiao-Yi Zhang, D. Towey, Tsong Y. Chen, Zheng Zheng, and Kai-Yuan Cai. Using partition information to prioritize test cases for fault localization. In *Proceedings of the IEEE 39th Annual International Computers, Software and Applications Conference*, COMPSAC'15, pages 121–126, 2015.
- Yunqian Zhang, Lin Chen, Bo Jiang, and Zhenyu Zhang. Wielding statistical fault localization statistically. In *Proceedings of the IEEE 23rd International Symposium on Software Reliability Engineering Workshops*, ISSREW'12, pages 189–194, 2012.
- Zhenyu Zhang, W. K. Chan, T. H. Tse, Peifeng Hu, and Xinming Wang. Is non-parametric hypothesis testing model robust for statistical fault localization? *Information and Software Technology*, 51(11):1573–1585, 2009a.
- Zhenyu Zhang, W. K. Chan, T. H. Tse, Bo Jiang, and Xinming Wang. Capturing propagation of infected program states. In *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the 10th ACM SIGSOFT Symposium on the Foundations of Software Engineering*, ESEC/FSE'09, pages 43–52, 2009b.
- Zhenyu Zhang, W. K. Chan, T. H. Tse, Yuen Tak Yu, and Peifeng Hu. Non-parametric statistical fault localization. *Journal of Systems and Software*, 84(6):885–905, 2011.
- Lei Zhao, Lina Wang, Zuoting Xiong, and Dongming Gao. Execution-aware fault localization based on the control flow analysis. In Rongbo Zhu, Yanchun Zhang, Baoxiang Liu, and Chunfeng Liu, editors, *Information Computing and Applications*, volume 6377 of *Lecture Notes in Computer Science*, pages 158–165. Springer, Berlin, 2010.
- Lei Zhao, Lina Wang, and Xiaodan Yin. Context-aware fault localization via control flow analysis. *Journal of Software*, 6(10):1977–1984, 2011a.
- Lei Zhao, Zhenyu Zhang, Lina Wang, and Xiaodan Yin. Paff: Fault localization via noise reduction on coverage vector. In *Proceedings of the 23rd International Conference on Software Engineering and Knowledge Engineering*, SEKE'11, pages 203–206, 2011b.
- Alice X. Zheng, Michael I. Jordan, Ben Liblit, Mayur Naik, and Alex Aiken. Statistical debugging: Simultaneous identification of multiple bugs. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML'06, pages 1105–1112, 2006.